

US009105068B2

(12) **United States Patent**
Lee et al.

(10) **Patent No.:** **US 9,105,068 B2**
(45) **Date of Patent:** **Aug. 11, 2015**

(54) **GRAMMAR MODEL FOR STRUCTURED SEARCH QUERIES**

USPC 707/769
See application file for complete search history.

(71) Applicants: **Yofay Kari Lee**, Palo Alto, CA (US);
Michael Benjamin Cohen, Menlo Park, CA (US); **Maxime Boucher**, Mountain View, CA (US); **Alisson Gusatti Azzolini**, San Francisco, CA (US); **Xiao Li**, Sunnyvale, CA (US); **Lars Eilstrup Rasmussen**, San Francisco, CA (US); **Kathryn Hymes**, Mountain View, CA (US); **Amy Campbell**, San Francisco, CA (US)

(56) **References Cited**

U.S. PATENT DOCUMENTS

| | | | |
|-----------|----|---------|--------------|
| 6,957,184 | B2 | 10/2005 | Schmid |
| 7,539,697 | B1 | 5/2009 | Akella |
| 7,752,326 | B2 | 7/2010 | Smit |
| 7,836,044 | B2 | 11/2010 | Kamvar |
| 7,933,843 | B1 | 4/2011 | Groll |
| 8,027,990 | B1 | 9/2011 | Mysen |
| 8,060,639 | B2 | 11/2011 | Smit |
| 8,112,529 | B2 | 2/2012 | Van Den Oord |
| 8,185,558 | B1 | 5/2012 | Narayanan |
| 8,239,364 | B2 | 8/2012 | Wable |
| 8,271,471 | B1 | 9/2012 | Kamvar |
| 8,271,546 | B2 | 9/2012 | Gibbs |
| 8,321,364 | B1 | 11/2012 | Gharpure |
| 8,364,709 | B1 | 1/2013 | Das |
| 8,386,465 | B2 | 2/2013 | Ansari |
| 8,407,200 | B2 | 3/2013 | Wable |
| 8,412,749 | B2 | 4/2013 | Fortuna |
| 8,538,960 | B2 | 9/2013 | Wong |
| 8,595,297 | B2 | 11/2013 | Marcucci |

(Continued)

OTHER PUBLICATIONS

(72) Inventors: **Yofay Kari Lee**, Palo Alto, CA (US);
Michael Benjamin Cohen, Menlo Park, CA (US); **Maxime Boucher**, Mountain View, CA (US); **Alisson Gusatti Azzolini**, San Francisco, CA (US); **Xiao Li**, Sunnyvale, CA (US); **Lars Eilstrup Rasmussen**, San Francisco, CA (US); **Kathryn Hymes**, Mountain View, CA (US); **Amy Campbell**, San Francisco, CA (US)

International Search Report and Written Opinion for International Application PCT/US2013/067214, Feb. 7, 2014.

(Continued)

(73) Assignee: **Facebook, Inc.**, Menlo Park, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 326 days.

Primary Examiner — Jacob F B  tit

Assistant Examiner — Kris Andersen

(74) *Attorney, Agent, or Firm* — Baker Botts L.L.P.

(21) Appl. No.: **13/674,695**

(22) Filed: **Nov. 12, 2012**

(65) **Prior Publication Data**

US 2014/0136564 A1 May 15, 2014

(51) **Int. Cl.**
G06F 17/30 (2006.01)
G06Q 50/00 (2012.01)

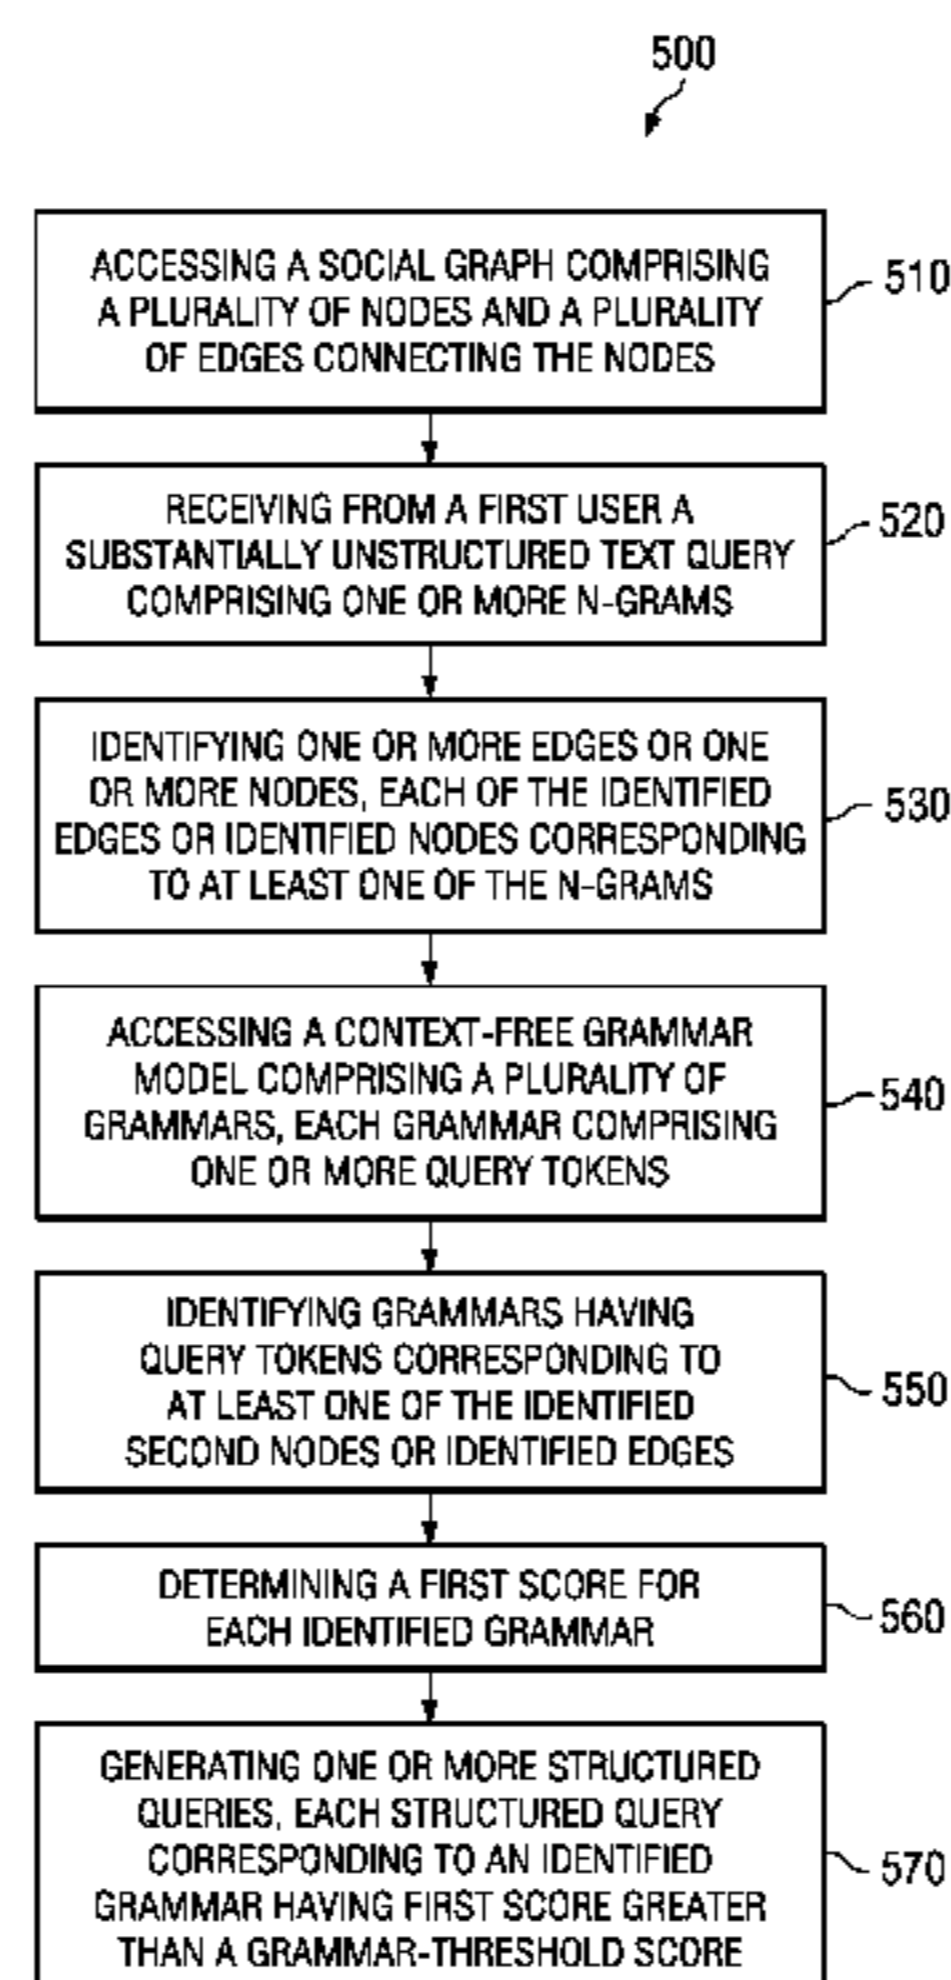
(52) **U.S. Cl.**
CPC **G06Q 50/01** (2013.01); **G06F 17/30867** (2013.01)

(58) **Field of Classification Search**
CPC G06F 17/30867; G06Q 50/10

(57) **ABSTRACT**

In one embodiment, a method includes accessing a social graph that includes a plurality of nodes and edges, receiving an unstructured text query, identifying nodes and edges that correspond to n-grams of the text query, accessing a context-free grammar model, identifying grammars having query tokens that correspond to the identified nodes and edges, determining a score for each identified grammar, and then generating structured queries based on the identified grammars based on strings generated by the grammars.

16 Claims, 5 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

8,601,027 B2 12/2013 Behforooz
 2005/0131872 A1 6/2005 Calbucci
 2007/0277100 A1 11/2007 Sheha
 2008/0126080 A1 5/2008 Saldanha
 2008/0270615 A1 10/2008 Centola
 2009/0006543 A1 1/2009 Smit
 2009/0228296 A1 9/2009 Ismalon
 2009/0281988 A1 11/2009 Yoo
 2010/0125604 A1 5/2010 Martinez
 2010/0235354 A1 9/2010 Gargaro
 2010/0281044 A1 11/2010 Gupta
 2011/0010367 A1 1/2011 Jockish
 2011/0078166 A1 3/2011 Oliver
 2011/0087534 A1 4/2011 Strebinger
 2011/0087968 A1 4/2011 Lakshmanan
 2011/0137932 A1 6/2011 Wable
 2011/0196855 A1 8/2011 Wable
 2011/0320470 A1* 12/2011 Williams et al. 707/767
 2012/0209832 A1* 8/2012 Neystadt et al. 707/723
 2012/0278127 A1 11/2012 Kirakosyan
 2012/0284329 A1 11/2012 Van Den Oord

2012/0311034 A1 12/2012 Goldband
 2013/0031106 A1 1/2013 Schechter
 2013/0041876 A1 2/2013 Dow
 2013/0066876 A1 3/2013 Raskino
 2013/0086057 A1 4/2013 Harrington
 2013/0198219 A1* 8/2013 Cohen et al. 707/760

OTHER PUBLICATIONS

European Patent Office Search Report for Application No. 13192539.
 8-1952, Feb. 24, 2014.
 "Customizing Results Snippets," Custom Search, Google Developers, <https://developers.google.com/custom-search/docs/snippets>, Nov. 21, 2012.
 Pynadath, David et al., "Generalized Queries on Probabilistic Context-Free Grammars," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, No. 1, Jan. 1998.
 Grune, Dick et al., Parsing Techniques, Chapter 13, 2008.
 U.S. Appl. No. 13/731,866, filed Dec. 31, 2012, Vee.
 U.S. Appl. No. 13/731,939, filed Dec. 31, 2012, Whitnah.
 U.S. Appl. No. 13/731,910, filed Dec. 31, 2012, Whitnah.
 U.S. Appl. No. 13/827,214, filed Mar. 14, 2013, Campbell.

* cited by examiner




FIG. 1




FIG. 6

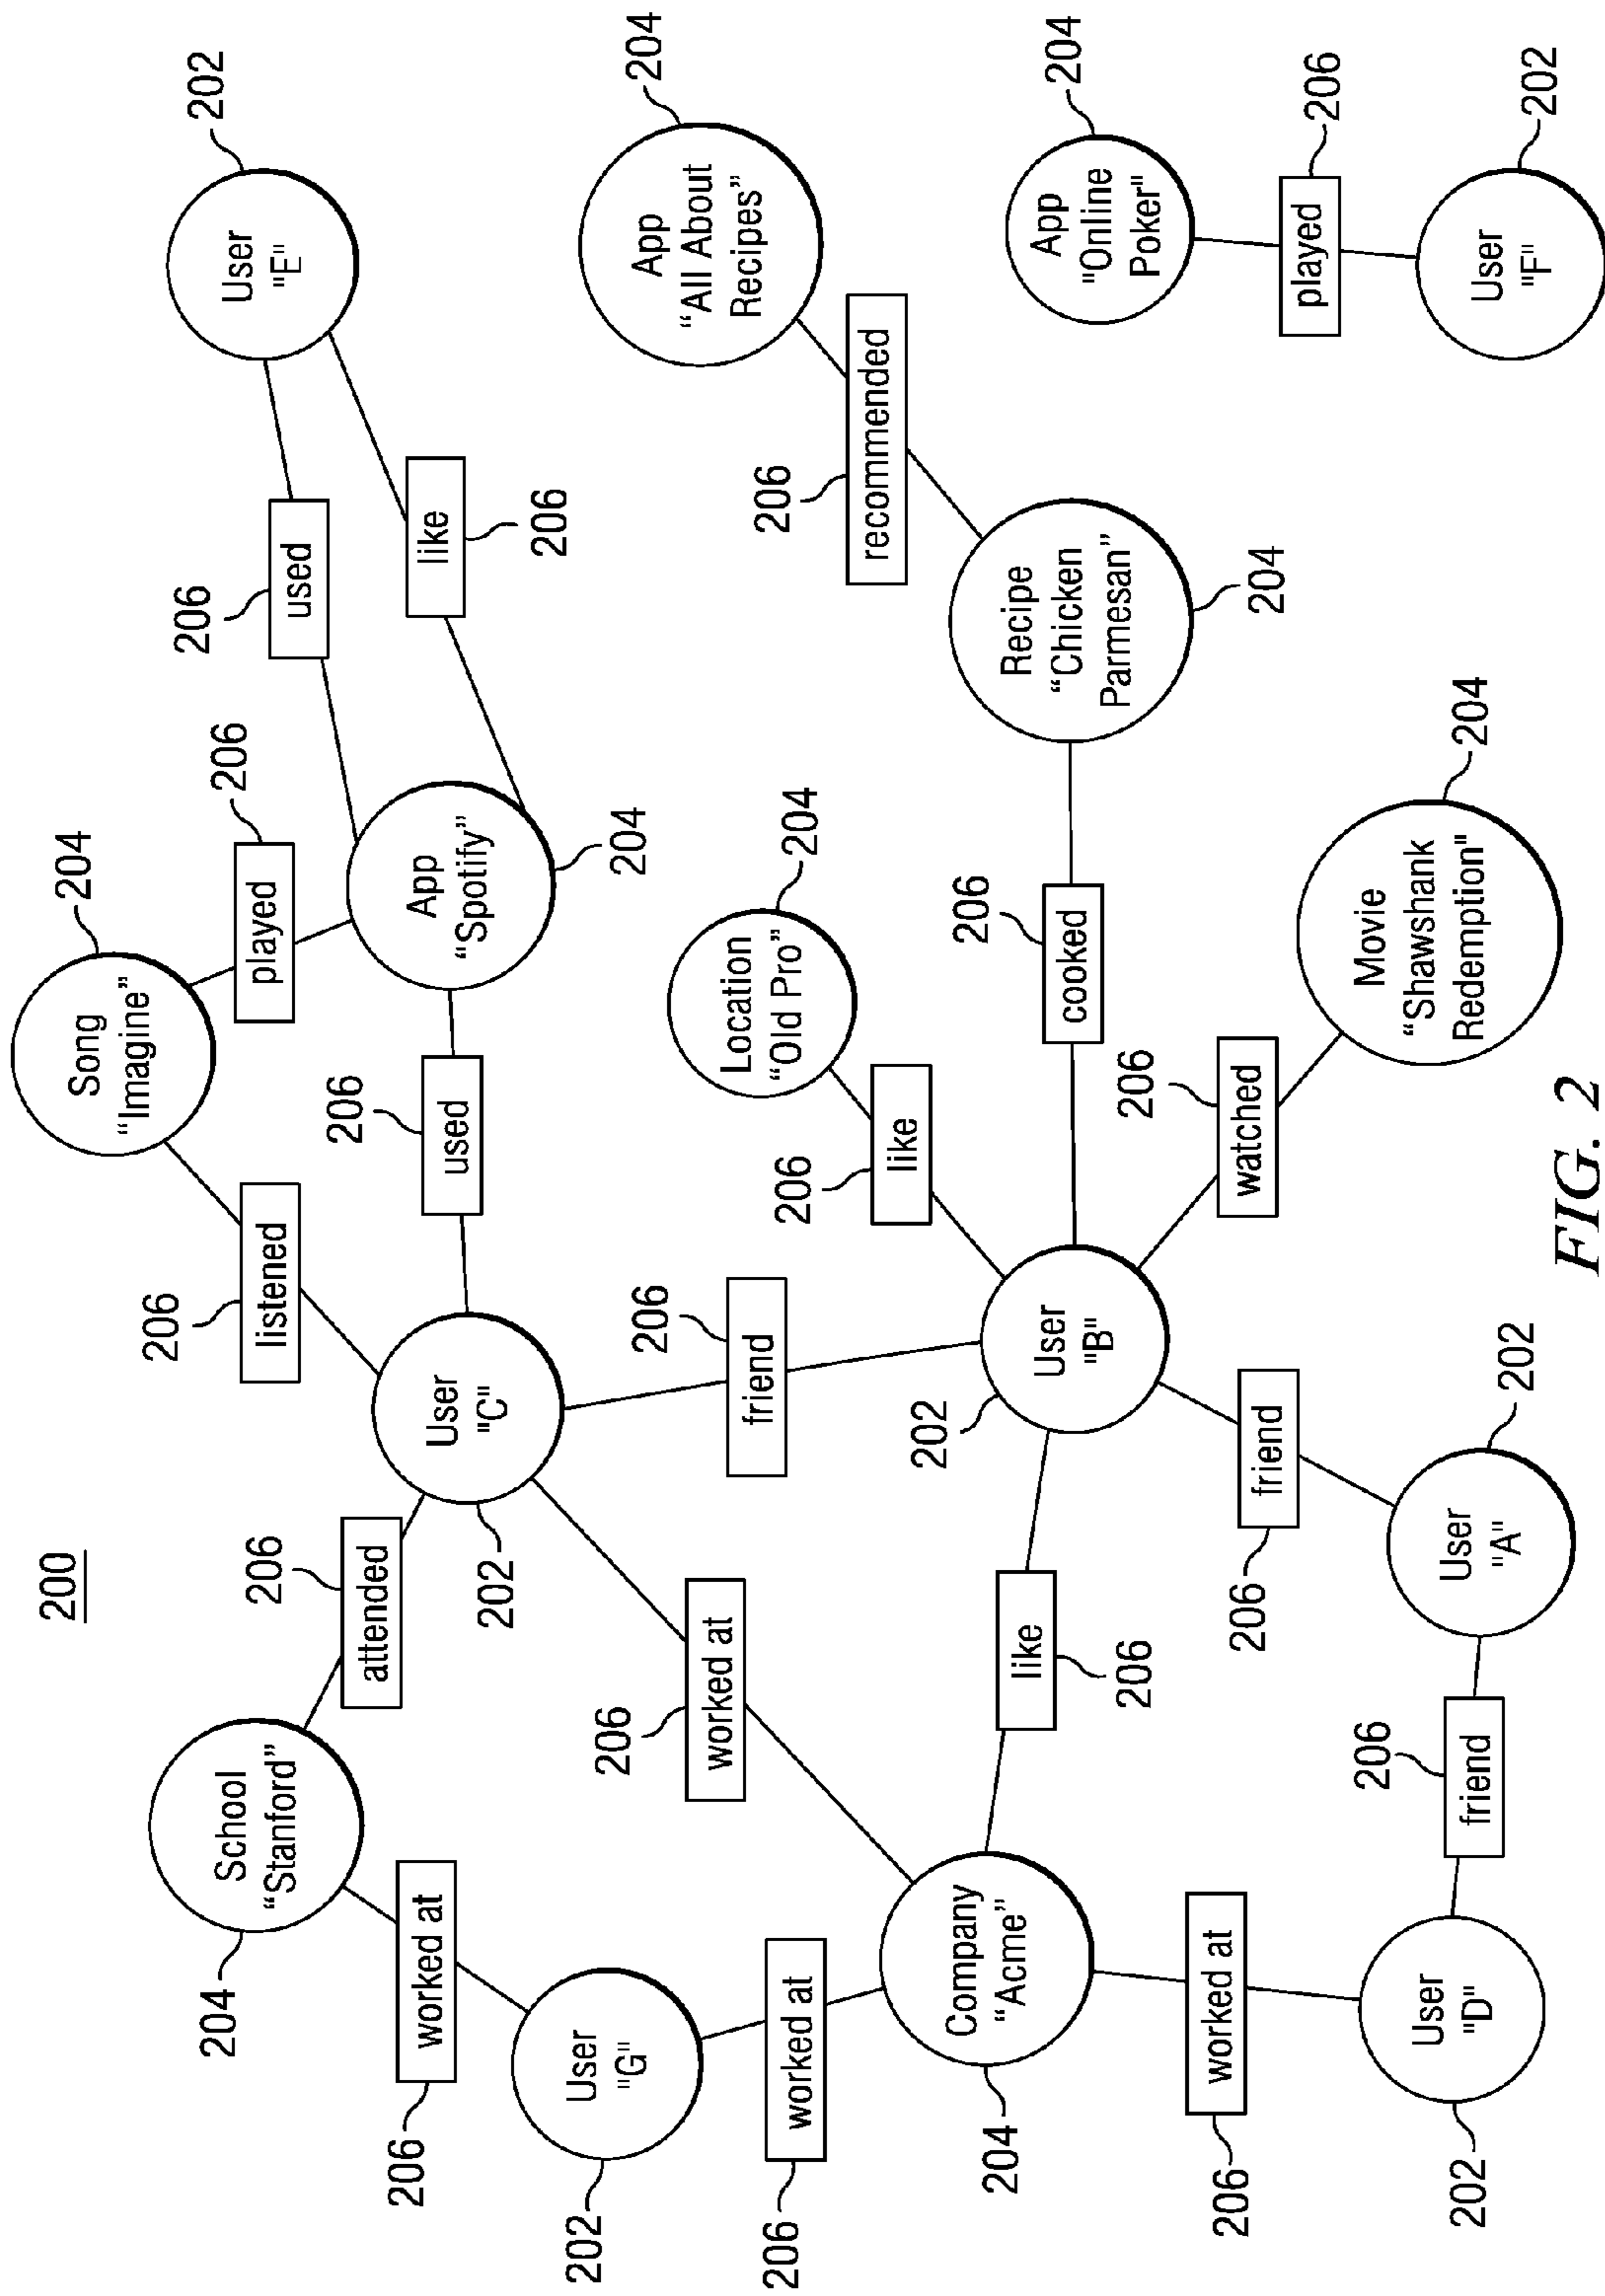


FIG. 2

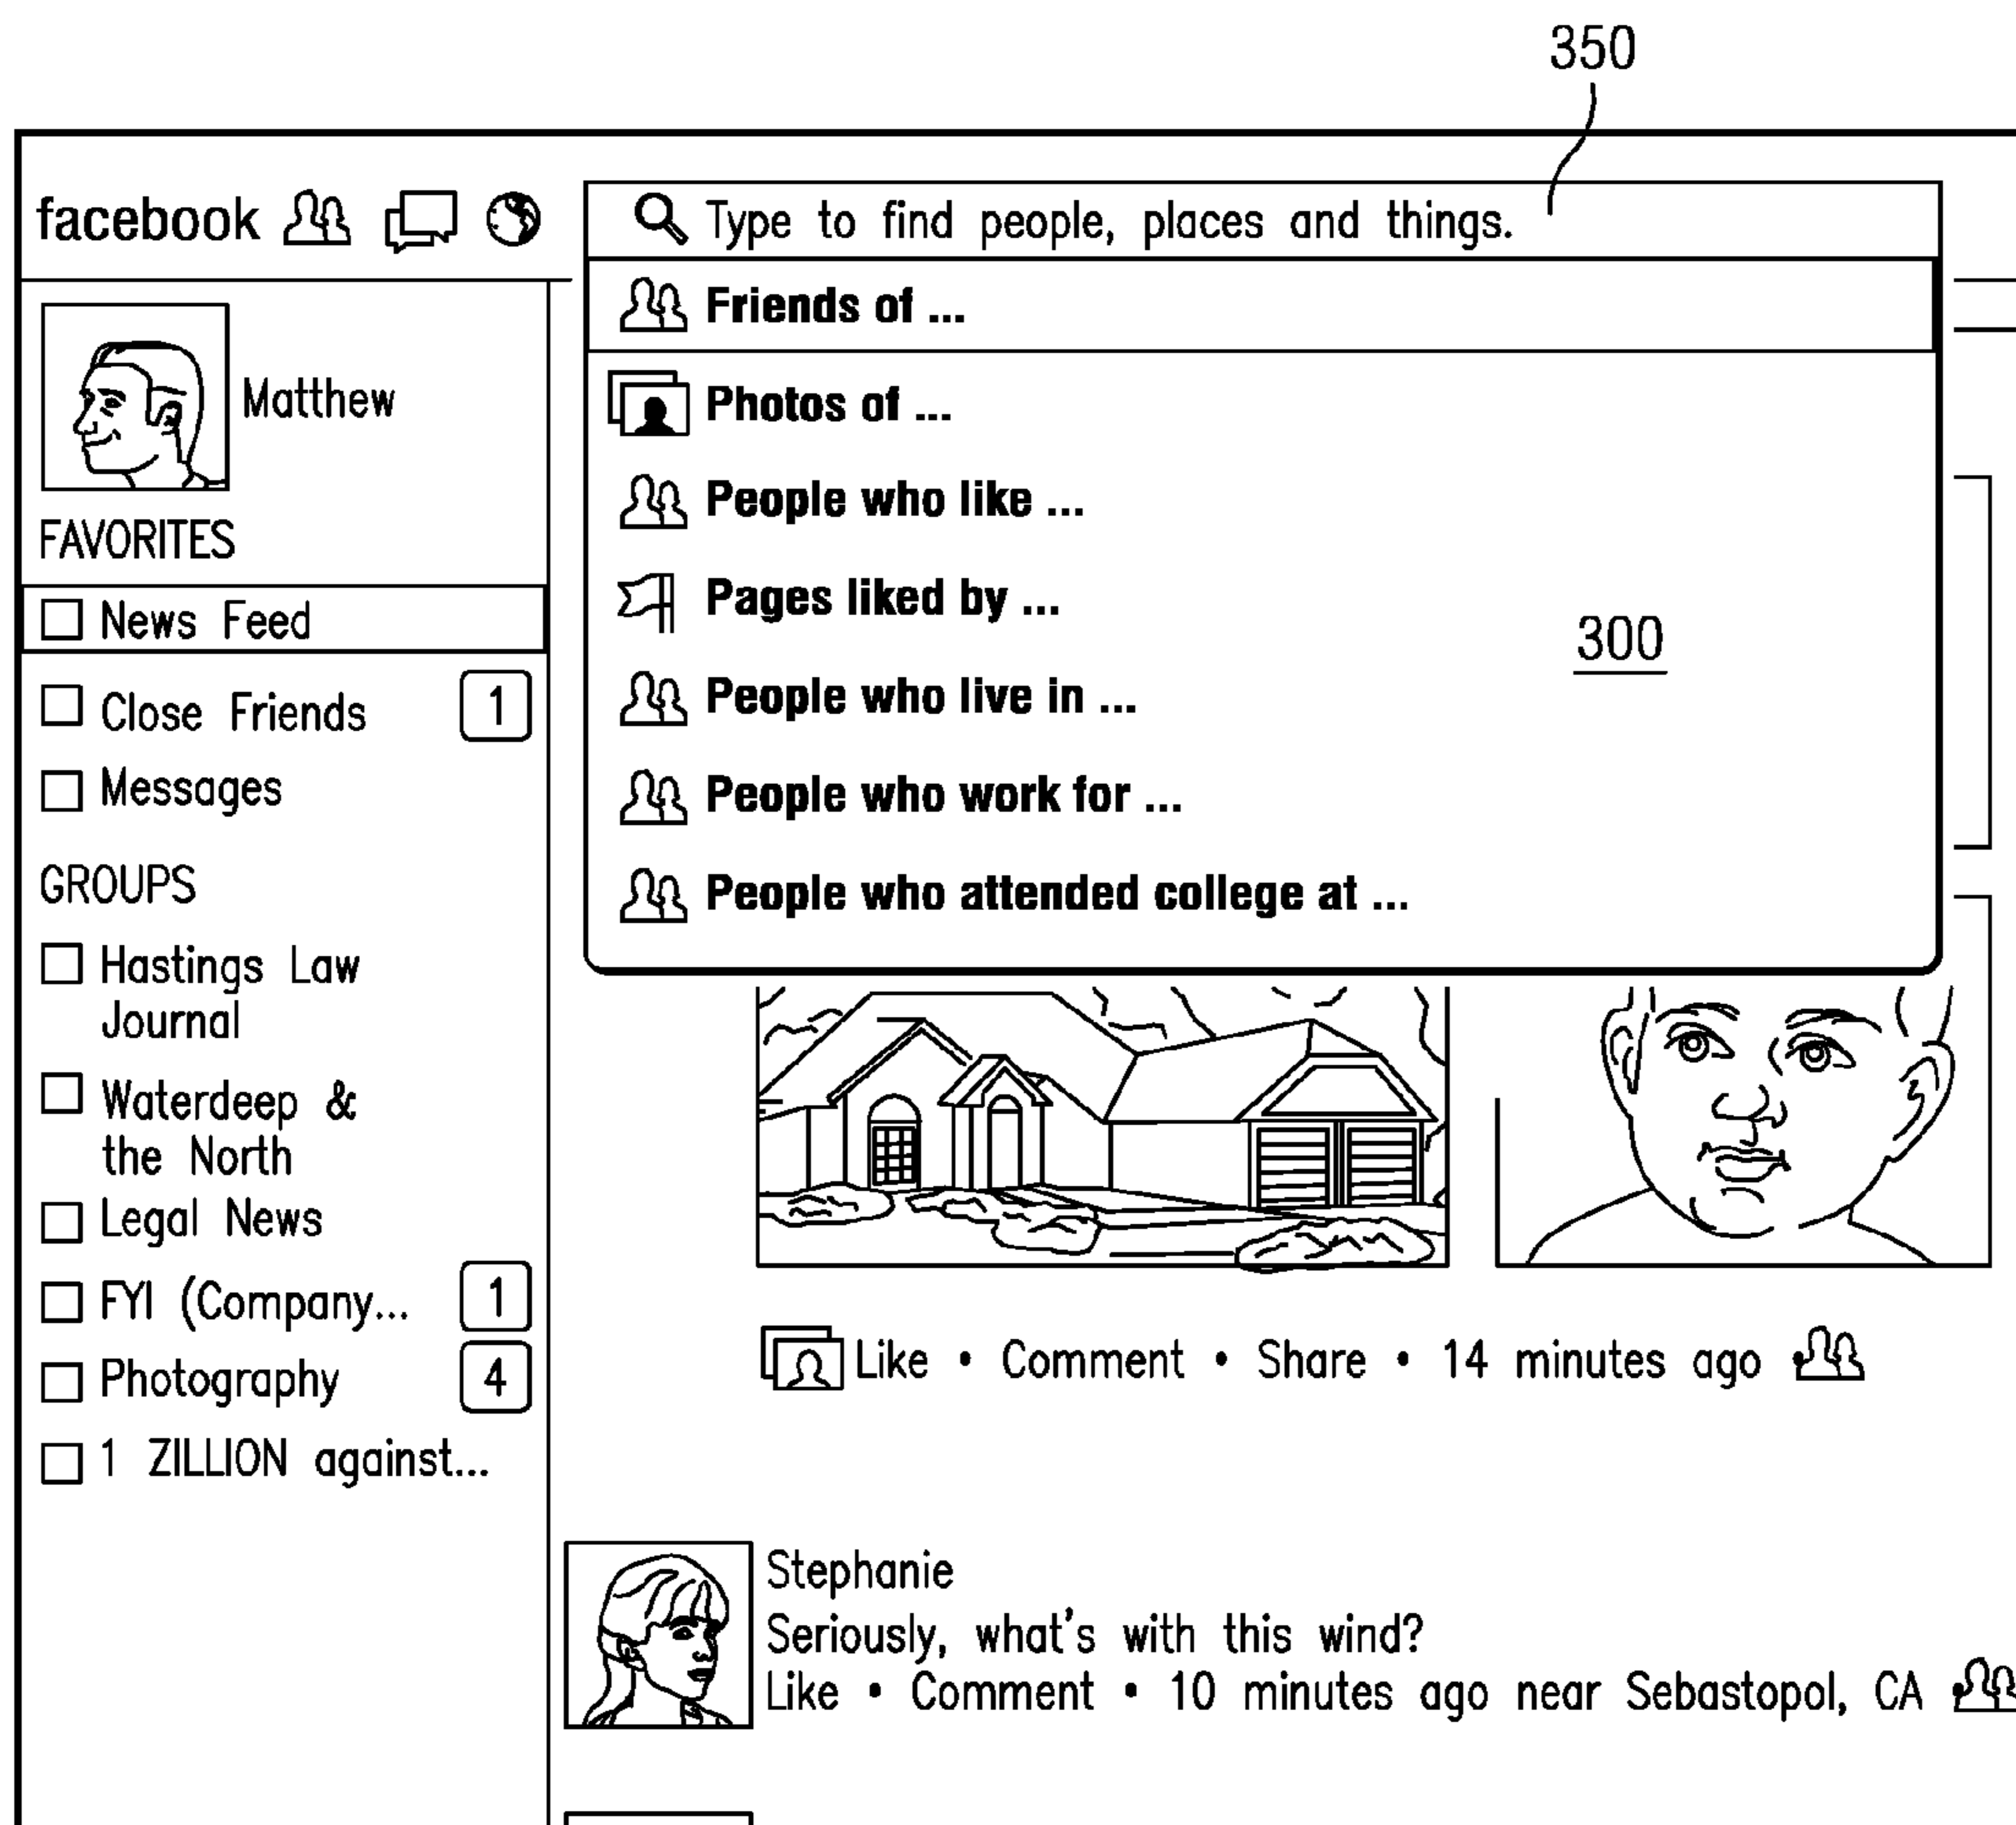


FIG. 3

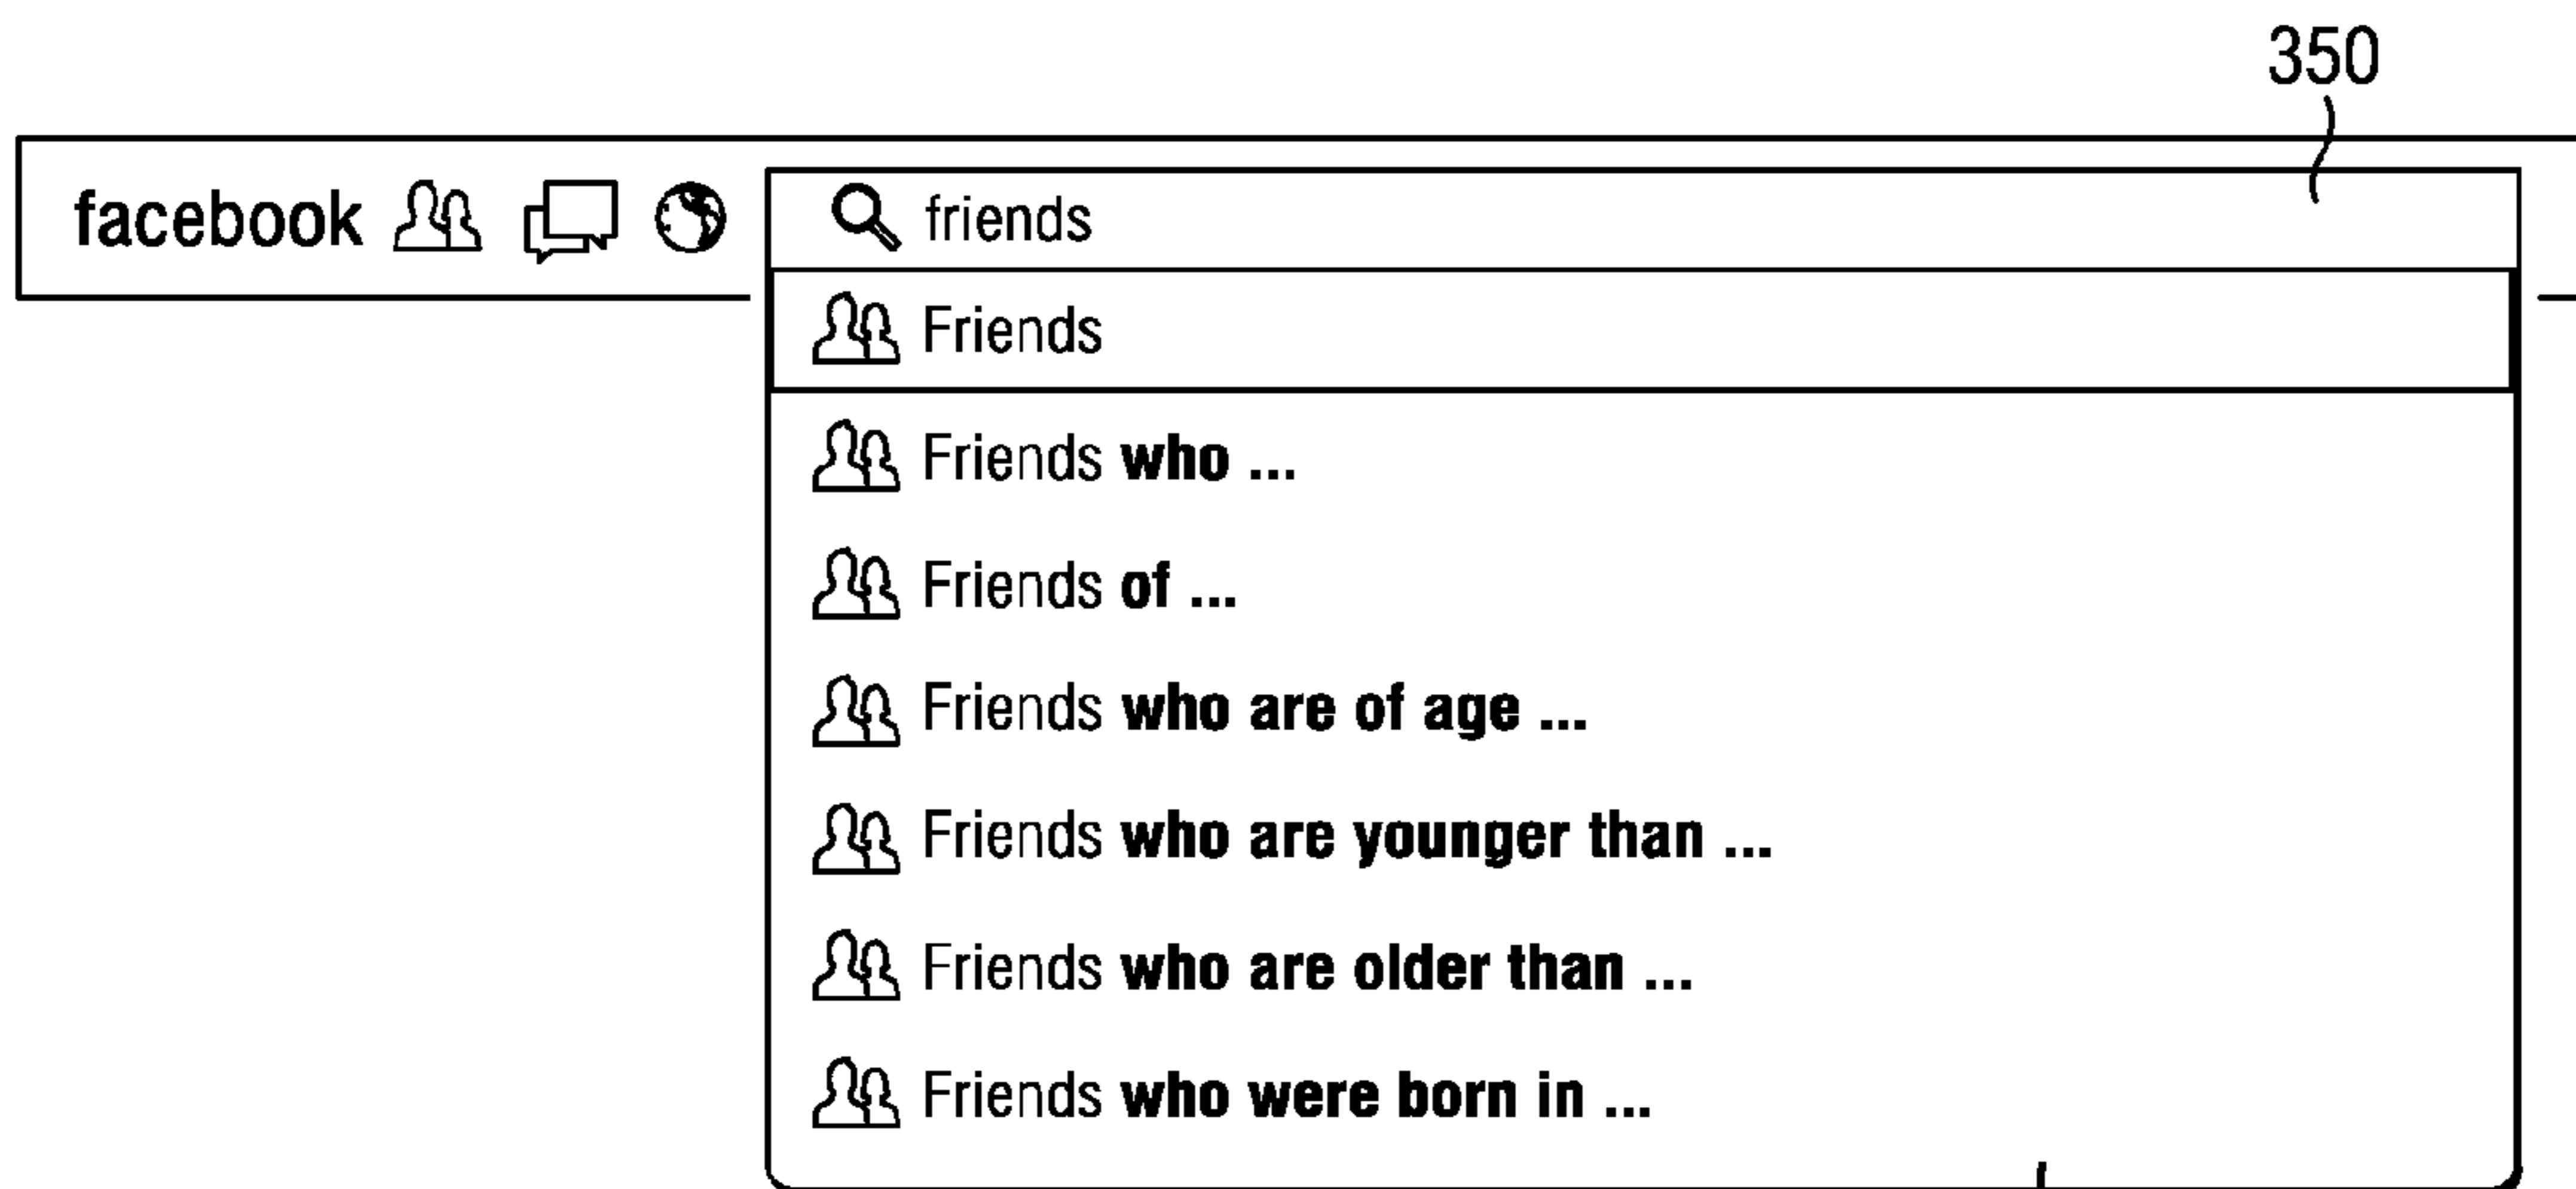


FIG. 4A

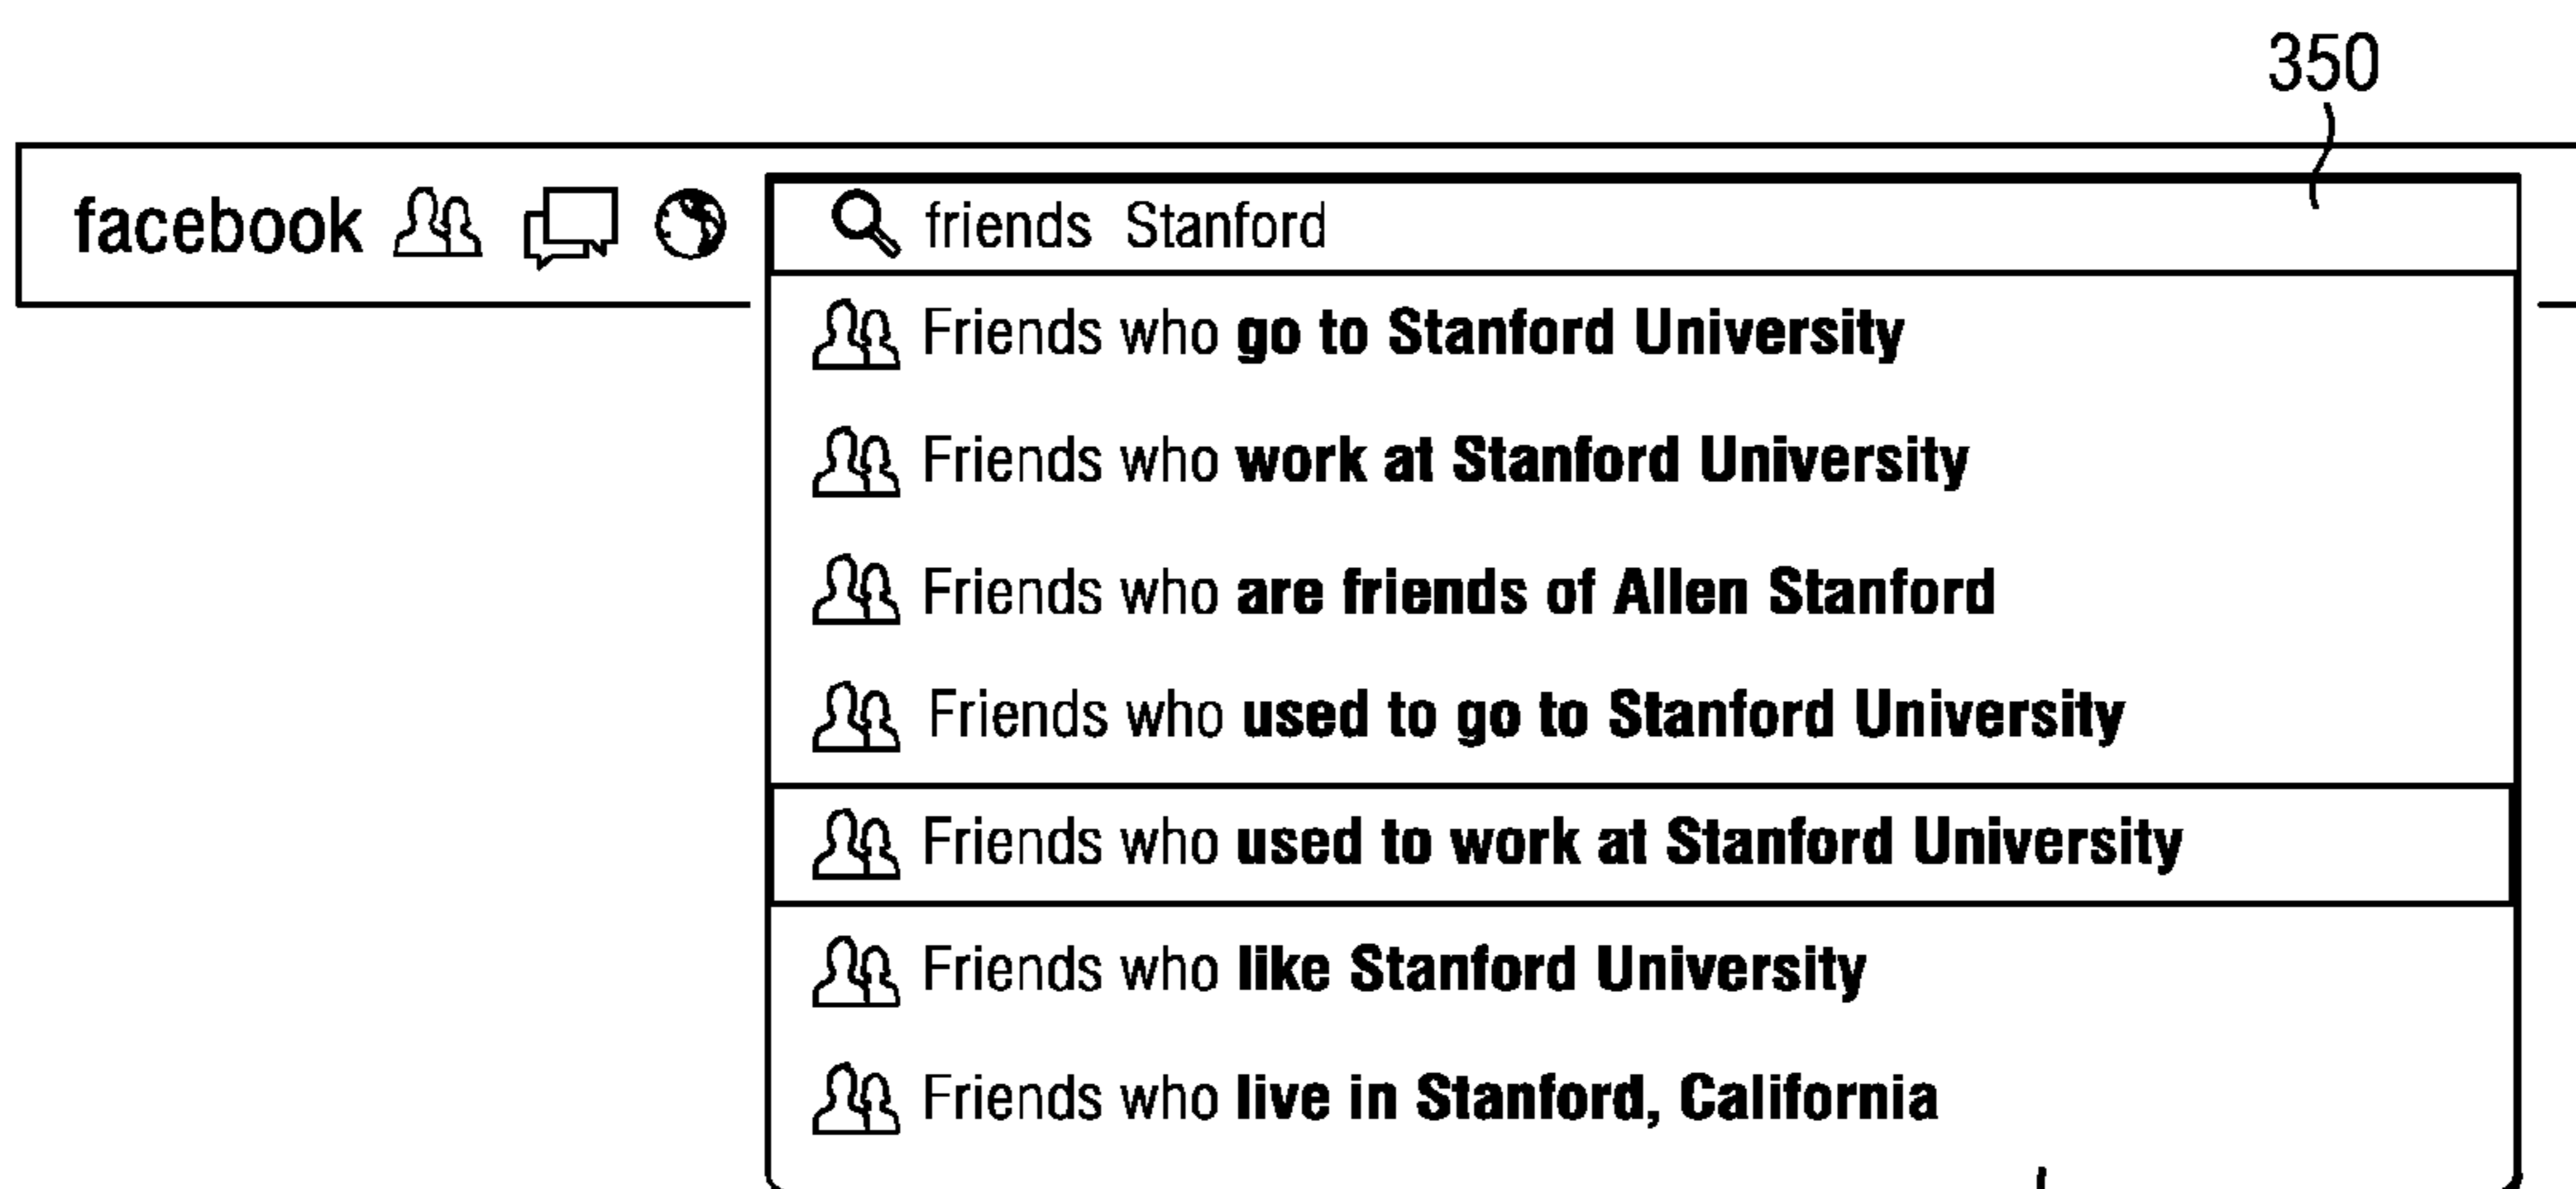


FIG. 4B

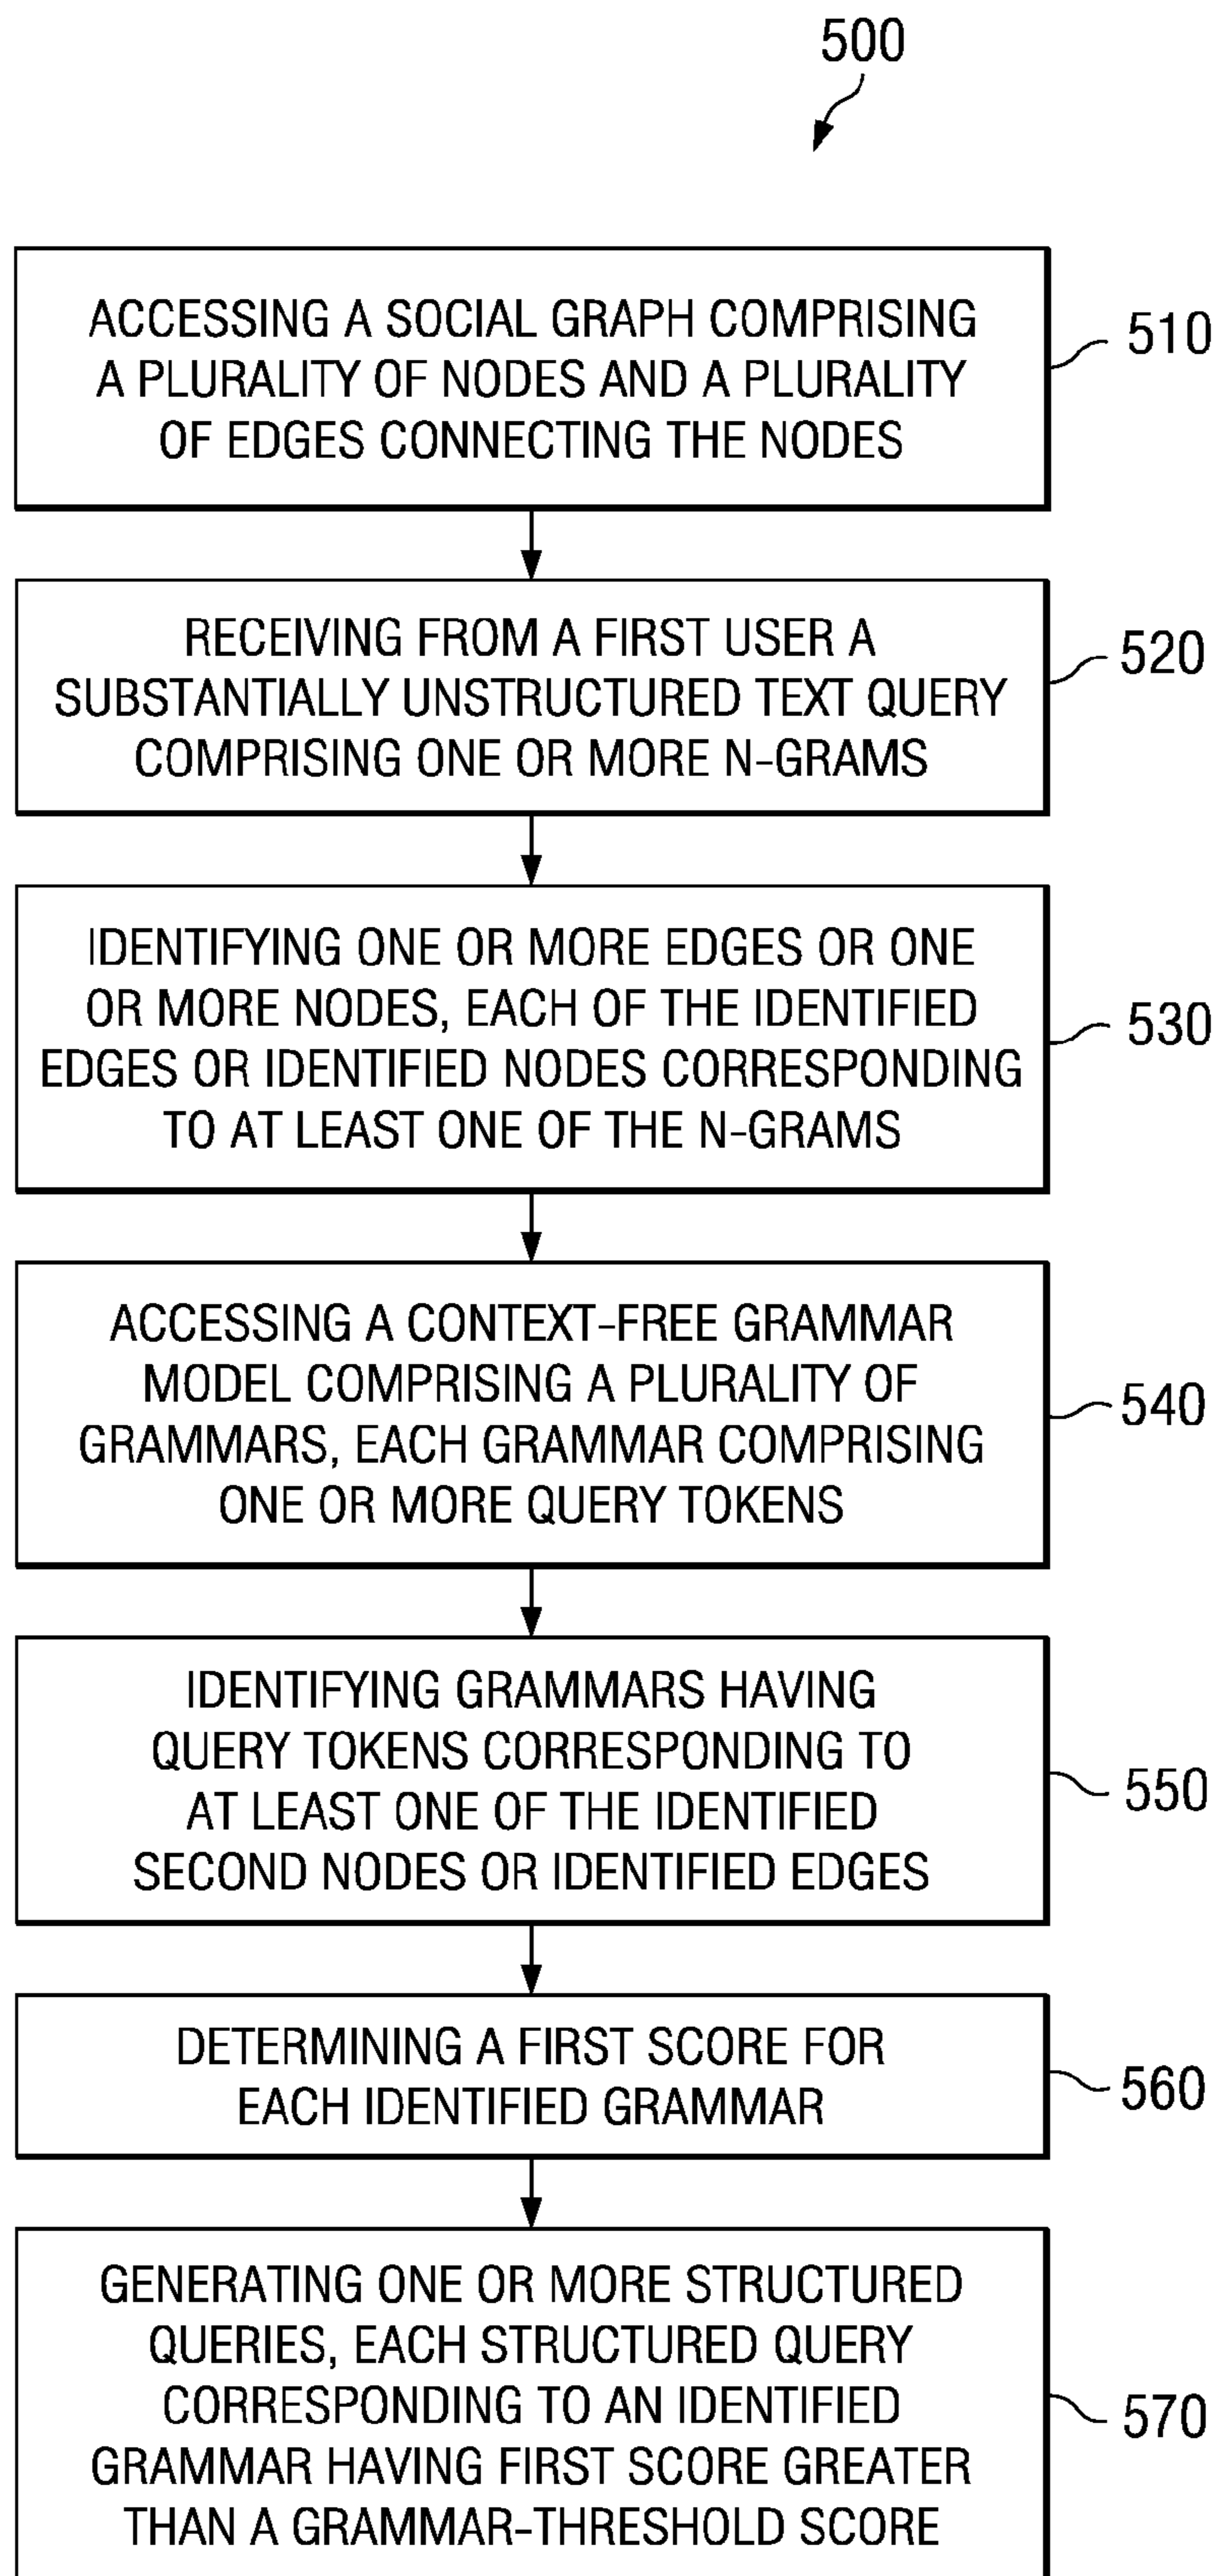


FIG. 5

1

GRAMMAR MODEL FOR STRUCTURED
SEARCH QUERIES

TECHNICAL FIELD

This disclosure generally relates to social graphs and performing searches for objects within a social-networking environment.

BACKGROUND

A social-networking system, which may include a social-networking website, may enable its users (such as persons or organizations) to interact with it and with each other through it. The social-networking system may, with input from a user, create and store in the social-networking system a user profile associated with the user. The user profile may include demographic information, communication-channel information, and information on personal interests of the user. The social-networking system may also, with input from a user, create and store a record of relationships of the user with other users of the social-networking system, as well as provide services (e.g. wall posts, photo-sharing, event organization, messaging, games, or advertisements) to facilitate social interaction between or among users.

The social-networking system may transmit over one or more networks content or messages related to its services to a mobile or other computing device of a user. A user may also install software applications on a mobile or other computing device of the user for accessing a user profile of the user and other data within the social-networking system. The social-networking system may generate a personalized set of content objects to display to a user, such as a newsfeed of aggregated stories of other users connected to the user.

Social-graph analysis views social relationships in terms of network theory consisting of nodes and edges. Nodes represent the individual actors within the networks, and edges represent the relationships between the actors. The resulting graph-based structures are often very complex. There can be many types of nodes and many types of edges for connecting nodes. In its simplest form, a social graph is a map of all of the relevant edges between all the nodes being studied.

SUMMARY OF PARTICULAR EMBODIMENTS

In particular embodiments, in response to a text query received from a user, a social-networking system may generate structured queries comprising query tokens that correspond to identified social-graph elements. By providing suggested structured queries in response to a user's text query, the social-networking system may provide a powerful way for users of an online social network to search for elements represented in a social graph based on their social-graph attributes and their relation to various social-graph elements.

In particular embodiments, the social-networking system may receive a substantially unstructured text query from a user. In response, the social-networking system may access a social graph and then parse the text query to identify social-graph elements that corresponded to n-grams from the text query. The social-networking system may identify these corresponding social-graph elements by determining a probability for each n-gram that it corresponds to a particular social-graph element. The social-networking system may then access a grammar model, such as a context-free grammar model. The identified social-graph elements may be used as terminal tokens ("query tokens") in the grammars of the grammar model, and each grammar may then be scored.

2

Grammars with a score greater than a threshold score may be used to generate structured queries that include query tokens referencing the identified social-graph elements. The structured queries may then be transmitted and displayed to the user, where the user can then select an appropriate query to search for the desired content.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates an example network environment associated with a social-networking system.

FIG. 2 illustrates an example social graph.

FIG. 3 illustrates an example webpage of an online social network.

FIGS. 4A-4B illustrate example queries of the social network.

FIG. 5 illustrates an example method for using a context-free grammar model to generate structured search queries.

FIG. 6 illustrates an example computer system.

DESCRIPTION OF EXAMPLE EMBODIMENTS

System Overview

FIG. 1 illustrates an example network environment 100 associated with a social-networking system. Network environment 100 includes a client system 130, a social-networking system 160, and a third-party system 170 connected to each other by a network 110. Although FIG. 1 illustrates a particular arrangement of client system 130, social-networking system 160, third-party system 170, and network 110, this disclosure contemplates any suitable arrangement of client system 130, social-networking system 160, third-party system 170, and network 110. As an example and not by way of limitation, two or more of client system 130, social-networking system 160, and third-party system 170 may be connected to each other directly, bypassing network 110. As another example, two or more of client system 130, social-networking system 160, and third-party system 170 may be physically or logically co-located with each other in whole or in part. Moreover, although FIG. 1 illustrates a particular number of client systems 130, social-networking systems 160, third-party systems 170, and networks 110, this disclosure contemplates any suitable number of client systems 130, social-networking systems 160, third-party systems 170, and networks 110. As an example and not by way of limitation, network environment 100 may include multiple client system 130, social-networking systems 160, third-party systems 170, and networks 110.

This disclosure contemplates any suitable network 110. As an example and not by way of limitation, one or more portions of network 110 may include an ad hoc network, an intranet, an extranet, a virtual private network (VPN), a local area network (LAN), a wireless LAN (WLAN), a wide area network (WAN), a wireless WAN (WWAN), a metropolitan area network (MAN), a portion of the Internet, a portion of the Public Switched Telephone Network (PSTN), a cellular telephone network, or a combination of two or more of these. Network 110 may include one or more networks 110.

Links 150 may connect client system 130, social-networking system 160, and third-party system 170 to communication network 110 or to each other. This disclosure contemplates any suitable links 150. In particular embodiments, one or more links 150 include one or more wireline (such as for example Digital Subscriber Line (DSL) or Data Over Cable Service Interface Specification (DOCSIS)), wireless (such as for example Wi-Fi or Worldwide Interoperability for Microwave Access (WiMAX)), or optical (such as for example

Synchronous Optical Network (SONET) or Synchronous Digital Hierarchy (SDH)) links. In particular embodiments, one or more links **150** each include an ad hoc network, an intranet, an extranet, a VPN, a LAN, a WLAN, a WAN, a WWAN, a MAN, a portion of the Internet, a portion of the PSTN, a cellular technology-based network, a satellite communications technology-based network, another link **150**, or a combination of two or more such links **150**. Links **150** need not necessarily be the same throughout network environment **100**. One or more first links **150** may differ in one or more respects from one or more second links **150**.

In particular embodiments, client system **130** may be an electronic device including hardware, software, or embedded logic components or a combination of two or more such components and capable of carrying out the appropriate functionalities implemented or supported by client system **130**. As an example and not by way of limitation, a client system **130** may include a computer system such as a desktop computer, notebook or laptop computer, netbook, a tablet computer, e-book reader, GPS device, camera, personal digital assistant (PDA), handheld electronic device, cellular telephone, smartphone, other suitable electronic device, or any suitable combination thereof. This disclosure contemplates any suitable client systems **130**. A client system **130** may enable a network user at client system **130** to access network **110**. A client system **130** may enable its user to communicate with other users at other client systems **130**.

In particular embodiments, client system **130** may include a web browser **132**, such as MICROSOFT INTERNET EXPLORER, GOOGLE CHROME or MOZILLA FIREFOX, and may have one or more add-ons, plug-ins, or other extensions, such as TOOLBAR or YAHOO TOOLBAR. A user at client system **130** may enter a Uniform Resource Locator (URL) or other address directing the web browser **132** to a particular server (such as server **162**, or a server associated with a third-party system **170**), and the web browser **132** may generate a Hyper Text Transfer Protocol (HTTP) request and communicate the HTTP request to server. The server may accept the HTTP request and communicate to client system **130** one or more Hyper Text Markup Language (HTML) files responsive to the HTTP request. Client system **130** may render a webpage based on the HTML files from the server for presentation to the user. This disclosure contemplates any suitable webpage files. As an example and not by way of limitation, webpages may render from HTML files, Extensible Hyper Text Markup Language (XHTML) files, or Extensible Markup Language (XML) files, according to particular needs. Such pages may also execute scripts such as, for example and without limitation, those written in JAVASCRIPT, JAVA, MICROSOFT SILVERLIGHT, combinations of markup language and scripts such as AJAX (Asynchronous JAVASCRIPT and XML), and the like. Herein, reference to a webpage encompasses one or more corresponding webpage files (which a browser may use to render the webpage) and vice versa, where appropriate.

In particular embodiments, social-networking system **160** may be a network-addressable computing system that can host an online social network. Social-networking system **160** may generate, store, receive, and transmit social-networking data, such as, for example, user-profile data, concept-profile data, social-graph information, or other suitable data related to the online social network. Social-networking system **160** may be accessed by the other components of network environment **100** either directly or via network **110**. In particular embodiments, social-networking system **160** may include one or more servers **162**. Each server **162** may be a unitary server or a distributed server spanning multiple computers or

multiple datacenters. Servers **162** may be of various types, such as, for example and without limitation, web server, news server, mail server, message server, advertising server, file server, application server, exchange server, database server, proxy server, another server suitable for performing functions or processes described herein, or any combination thereof. In particular embodiments, each server **162** may include hardware, software, or embedded logic components or a combination of two or more such components for carrying out the appropriate functionalities implemented or supported by server **162**. In particular embodiments, social-networking system **164** may include one or more data stores **164**. Data stores **164** may be used to store various types of information. In particular embodiments, the information stored in data stores **164** may be organized according to specific data structures. In particular embodiments, each data store **164** may be a relational database. Particular embodiments may provide interfaces that enable a client system **130**, a social-networking system **160**, or a third-party system **170** to manage, retrieve, modify, add, or delete, the information stored in data store **164**.

In particular embodiments, social-networking system **160** may store one or more social graphs in one or more data stores **164**. In particular embodiments, a social graph may include multiple nodes—which may include multiple user nodes (each corresponding to a particular user) or multiple concept nodes (each corresponding to a particular concept)—and multiple edges connecting the nodes. Social-networking system **160** may provide users of the online social network the ability to communicate and interact with other users. In particular embodiments, users may join the online social network via social-networking system **160** and then add connections (i.e., relationships) to a number of other users of social-networking system **160** whom they want to be connected to. Herein, the term “friend” may refer to any other user of social-networking system **160** with whom a user has formed a connection, association, or relationship via social-networking system **160**.

In particular embodiments, social-networking system **160** may provide users with the ability to take actions on various types of items or objects, supported by social-networking system **160**. As an example and not by way of limitation, the items and objects may include groups or social networks to which users of social-networking system **160** may belong, events or calendar entries in which a user might be interested, computer-based applications that a user may use, transactions that allow users to buy or sell items via the service, interactions with advertisements that a user may perform, or other suitable items or objects. A user may interact with anything that is capable of being represented in social-networking system **160** or by an external system of third-party system **170**, which is separate from social-networking system **160** and coupled to social-networking system **160** via a network **110**.

In particular embodiments, social-networking system **160** may be capable of linking a variety of entities. As an example and not by way of limitation, social-networking system **160** may enable users to interact with each other as well as receive content from third-party systems **170** or other entities, or to allow users to interact with these entities through an application programming interfaces (API) or other communication channels.

In particular embodiments, a third-party system **170** may include one or more types of servers, one or more data stores, one or more interfaces, including but not limited to APIs, one or more web services, one or more content sources, one or more networks, or any other suitable components, e.g., that

servers may communicate with. A third-party system **170** may be operated by a different entity from an entity operating social-networking system **160**. In particular embodiments, however, social-networking system **160** and third-party systems **170** may operate in conjunction with each other to provide social-networking services to users of social-networking system **160** or third-party systems **170**. In this sense, social-networking system **160** may provide a platform, or backbone, which other systems, such as third-party systems **170**, may use to provide social-networking services and functionality to users across the Internet.

In particular embodiments, a third-party system **170** may include a third-party content object provider. A third-party content object provider may include one or more sources of content objects, which may be communicated to a client system **130**. As an example and not by way of limitation, content objects may include information regarding things or activities of interest to the user, such as, for example, movie show times, movie reviews, restaurant reviews, restaurant menus, product information and reviews, or other suitable information. As another example and not by way of limitation, content objects may include incentive content objects, such as coupons, discount tickets, gift certificates, or other suitable incentive objects.

In particular embodiments, social-networking system **160** also includes user-generated content objects, which may enhance a user's interactions with social-networking system **160**. User-generated content may include anything a user can add, upload, send, or "post" to social-networking system **160**. As an example and not by way of limitation, a user communicates posts to social-networking system **160** from a client system **130**. Posts may include data such as status updates or other textual data, location information, photos, videos, links, music or other similar data or media. Content may also be added to social-networking system **160** by a third-party through a "communication channel," such as a newsfeed or stream.

In particular embodiments, social-networking system **160** may include a variety of servers, sub-systems, programs, modules, logs, and data stores. In particular embodiments, social-networking system **160** may include one or more of the following: a web server, action logger, API-request server, relevance-and-ranking engine, content-object classifier, notification controller, action log, third-party-content-object-exposure log, inference module, authorization/privacy server, search module, ad-targeting module, user-interface module, user-profile store, connection store, third-party content store, or location store. Social-networking system **160** may also include suitable components such as network interfaces, security mechanisms, load balancers, failover servers, management-and-network-operations consoles, other suitable components, or any suitable combination thereof. In particular embodiments, social-networking system **160** may include one or more user-profile stores for storing user profiles. A user profile may include, for example, biographic information, demographic information, behavioral information, social information, or other types of descriptive information, such as work experience, educational history, hobbies or preferences, interests, affinities, or location. Interest information may include interests related to one or more categories. Categories may be general or specific. As an example and not by way of limitation, if a user "likes" an article about a brand of shoes the category may be the brand, or the general category of "shoes" or "clothing." A connection store may be used for storing connection information about users. The connection information may indicate users who have similar or common work experience, group memberships, hobbies, educational

history, or are in any way related or share common attributes. The connection information may also include user-defined connections between different users and content (both internal and external). A web server may be used for linking social-networking system **160** to one or more client systems **130** or one or more third-party system **170** via network **110**. The web server may include a mail server or other messaging functionality for receiving and routing messages between social-networking system **160** and one or more client systems **130**. An API-request server may allow a third-party system **170** to access information from social-networking system **160** by calling one or more APIs. An action logger may be used to receive communications from a web server about a user's actions on or off social-networking system **160**. In conjunction with the action log, a third-party-content-object log may be maintained of user exposures to third-party-content objects. A notification controller may provide information regarding content objects to a client system **130**. Information may be pushed to a client system **130** as notifications, or information may be pulled from client system **130** responsive to a request received from client system **130**. Authorization servers may be used to enforce one or more privacy settings of the users of social-networking system **160**. A privacy setting of a user determines how particular information associated with a user can be shared. The authorization server may allow users to opt in or opt out of having their actions logged by social-networking system **160** or shared with other systems (e.g., third-party system **170**), such as, for example, by setting appropriate privacy settings. Third-party-content-object stores may be used to store content objects received from third parties, such as a third-party system **170**. Location stores may be used for storing location information received from client systems **130** associated with users. Ad-pricing modules may combine social information, the current time, location information, or other suitable information to provide relevant advertisements, in the form of notifications, to a user.

Social Graphs

FIG. 2 illustrates example social graph **200**. In particular embodiments, social-networking system **160** may store one or more social graphs **200** in one or more data stores. In particular embodiments, social graph **200** may include multiple nodes—which may include multiple user nodes **202** or multiple concept nodes **204**—and multiple edges **206** connecting the nodes. Example social graph **200** illustrated in FIG. 2 is shown, for didactic purposes, in a two-dimensional visual map representation. In particular embodiments, a social-networking system **160**, client system **130**, or third-party system **170** may access social graph **200** and related social-graph information for suitable applications. The nodes and edges of social graph **200** may be stored as data objects, for example, in a data store (such as a social-graph database). Such a data store may include one or more searchable or queryable indexes of nodes or edges of social graph **200**.

In particular embodiments, a user node **202** may correspond to a user of social-networking system **160**. As an example and not by way of limitation, a user may be an individual (human user), an entity (e.g., an enterprise, business, or third-party application), or a group (e.g., of individuals or entities) that interacts or communicates with or over social-networking system **160**. In particular embodiments, when a user registers for an account with social-networking system **160**, social-networking system **160** may create a user node **202** corresponding to the user, and store the user node **202** in one or more data stores. Users and user nodes **202** described herein may, where appropriate, refer to registered users and user nodes **202** associated with registered users. In addition or as an alternative, users and user nodes **202**

described herein may, where appropriate, refer to users that have not registered with social-networking system 160. In particular embodiments, a user node 202 may be associated with information provided by a user or information gathered by various systems, including social-networking system 160. As an example and not by way of limitation, a user may provide his or her name, profile picture, contact information, birth date, sex, marital status, family status, employment, education background, preferences, interests, or other demographic information. In particular embodiments, a user node 202 may be associated with one or more data objects corresponding to information associated with a user. In particular embodiments, a user node 202 may correspond to one or more webpages.

In particular embodiments, a concept node 204 may correspond to a concept. As an example and not by way of limitation, a concept may correspond to a place (such as, for example, a movie theater, restaurant, landmark, or city); a website (such as, for example, a website associated with social-network system 160 or a third-party website associated with a web-application server); an entity (such as, for example, a person, business, group, sports team, or celebrity); a resource (such as, for example, an audio file, video file, digital photo, text file, structured document, or application) which may be located within social-networking system 160 or on an external server, such as a web-application server; real or intellectual property (such as, for example, a sculpture, painting, movie, game, song, idea, photograph, or written work); a game; an activity; an idea or theory; another suitable concept; or two or more such concepts. A concept node 204 may be associated with information of a concept provided by a user or information gathered by various systems, including social-networking system 160. As an example and not by way of limitation, information of a concept may include a name or a title; one or more images (e.g., an image of the cover page of a book); a location (e.g., an address or a geographical location); a website (which may be associated with a URL); contact information (e.g., a phone number or an email address); other suitable concept information; or any suitable combination of such information. In particular embodiments, a concept node 204 may be associated with one or more data objects corresponding to information associated with concept node 204. In particular embodiments, a concept node 204 may correspond to one or more webpages.

In particular embodiments, a node in social graph 200 may represent or be represented by a webpage (which may be referred to as a “profile page”). Profile pages may be hosted by or accessible to social-networking system 160. Profile pages may also be hosted on third-party websites associated with a third-party server 170. As an example and not by way of limitation, a profile page corresponding to a particular external webpage may be the particular external webpage and the profile page may correspond to a particular concept node 204. Profile pages may be viewable by all or a selected subset of other users. As an example and not by way of limitation, a user node 202 may have a corresponding user-profile page in which the corresponding user may add content, make declarations, or otherwise express himself or herself. As another example and not by way of limitation, a concept node 204 may have a corresponding concept-profile page in which one or more users may add content, make declarations, or express themselves, particularly in relation to the concept corresponding to concept node 204.

In particular embodiments, a concept node 204 may represent a third-party webpage or resource hosted by a third-party system 170. The third-party webpage or resource may include, among other elements, content, a selectable or other

icon, or other inter-actable object (which may be implemented, for example, in JavaScript, AJAX, or PHP codes) representing an action or activity. As an example and not by way of limitation, a third-party webpage may include a selectable icon such as “like,” “check in,” “eat,” “recommend,” or another suitable action or activity. A user viewing the third-party webpage may perform an action by selecting one of the icons (e.g., “eat”), causing a client system 130 to transmit to social-networking system 160 a message indicating the user’s action. In response to the message, social-networking system 160 may create an edge (e.g., an “eat” edge) between a user node 202 corresponding to the user and a concept node 204 corresponding to the third-party webpage or resource and store edge 206 in one or more data stores.

In particular embodiments, a pair of nodes in social graph 200 may be connected to each other by one or more edges 206. An edge 206 connecting a pair of nodes may represent a relationship between the pair of nodes. In particular embodiments, an edge 206 may include or represent one or more data objects or attributes corresponding to the relationship between a pair of nodes. As an example and not by way of limitation, a first user may indicate that a second user is a “friend” of the first user. In response to this indication, social-networking system 160 may transmit a “friend request” to the second user. If the second user confirms the “friend request,” social-networking system 160 may create an edge 206 connecting the first user’s user node 202 to the second user’s user node 202 in social graph 200 and store edge 206 as social-graph information in one or more of data stores 24. In the example of FIG. 2, social graph 200 includes an edge 206 indicating a friend relation between user nodes 202 of user “A” and user “B” and an edge indicating a friend relation between user nodes 202 of user “C” and user “B.” Although this disclosure describes or illustrates particular edges 206 with particular attributes connecting particular user nodes 202, this disclosure contemplates any suitable edges 206 with any suitable attributes connecting user nodes 202. As an example and not by way of limitation, an edge 206 may represent a friendship, family relationship, business or employment relationship, fan relationship, follower relationship, visitor relationship, subscriber relationship, superior/subordinate relationship, reciprocal relationship, non-reciprocal relationship, another suitable type of relationship, or two or more such relationships. Moreover, although this disclosure generally describes nodes as being connected, this disclosure also describes users or concepts as being connected. Herein, references to users or concepts being connected may, where appropriate, refer to the nodes corresponding to those users or concepts being connected in social graph 200 by one or more edges 206.

In particular embodiments, an edge 206 between a user node 202 and a concept node 204 may represent a particular action or activity performed by a user associated with user node 202 toward a concept associated with a concept node 204. As an example and not by way of limitation, as illustrated in FIG. 2, a user may “like,” “attended,” “played,” “listened,” “cooked,” “worked at,” or “watched” a concept, each of which may correspond to an edge type or subtype. A concept-profile page corresponding to a concept node 204 may include, for example, a selectable “check in” icon (such as, for example, a clickable “check in” icon) or a selectable “add to favorites” icon. Similarly, after a user clicks these icons, social-networking system 160 may create a “favorite” edge or a “check in” edge in response to a user’s action corresponding to a respective action. As another example and not by way of limitation, a user (user “C”) may listen to a particular song (“Imagine”) using a particular application (SPOTIFY, which

is an online music application). In this case, social-networking system 160 may create a “listened” edge 206 and a “used” edge (as illustrated in FIG. 2) between user nodes 202 corresponding to the user and concept nodes 204 corresponding to the song and application to indicate that the user listened to the song and used the application. Moreover, social-networking system 160 may create a “played” edge 206 (as illustrated in FIG. 2) between concept nodes 204 corresponding to the song and the application to indicate that the particular song was played by the particular application. In this case, “played” edge 206 corresponds to an action performed by an external application (SPOTIFY) on an external audio file (the song “Imagine”). Although this disclosure describes particular edges 206 with particular attributes connecting user nodes 202 and concept nodes 204, this disclosure contemplates any suitable edges 206 with any suitable attributes connecting user nodes 202 and concept nodes 204. Moreover, although this disclosure describes edges between a user node 202 and a concept node 204 representing a single relationship, this disclosure contemplates edges between a user node 202 and a concept node 204 representing one or more relationships. As an example and not by way of limitation, an edge 206 may represent both that a user likes and has used at a particular concept. Alternatively, another edge 206 may represent each type of relationship (or multiples of a single relationship) between a user node 202 and a concept node 204 (as illustrated in FIG. 2 between user node 202 for user “E” and concept node 204 for “SPOTIFY”).

In particular embodiments, social-networking system 160 may create an edge 206 between a user node 202 and a concept node 204 in social graph 200. As an example and not by way of limitation, a user viewing a concept-profile page (such as, for example, by using a web browser or a special-purpose application hosted by the user’s client system 130) may indicate that he or she likes the concept represented by the concept node 204 by clicking or selecting a “Like” icon, which may cause the user’s client system 130 to transmit to social-networking system 160 a message indicating the user’s liking of the concept associated with the concept-profile page. In response to the message, social-networking system 160 may create an edge 206 between user node 202 associated with the user and concept node 204, as illustrated by “like” edge 206 between the user and concept node 204. In particular embodiments, social-networking system 160 may store an edge 206 in one or more data stores. In particular embodiments, an edge 206 may be automatically formed by social-networking system 160 in response to a particular user action. As an example and not by way of limitation, if a first user uploads a picture, watches a movie, or listens to a song, an edge 206 may be formed between user node 202 corresponding to the first user and concept nodes 204 corresponding to those concepts. Although this disclosure describes forming particular edges 206 in particular manners, this disclosure contemplates forming any suitable edges 206 in any suitable manner.

Typeahead Processes

In particular embodiments, one or more client-side and/or backend (server-side) processes implement and utilize a “typeahead” feature to automatically attempt to match concepts corresponding to respective existing user nodes 202 or concept nodes 204 to information currently being entered by a user in an input form rendered in conjunction with a requested webpage, such as a user-profile page, which may be hosted or accessible in, by the social-networking system 160. In particular embodiments, as a user is entering text to make a declaration, the typeahead feature attempts to match the string of textual characters being entered in the declaration to

strings of characters (e.g., names) corresponding to existing concepts (or users) and corresponding concept (or user) nodes in the social graph 200. In particular embodiments, when a match is found, the typeahead feature may automatically populate the form with a reference to the node (such as, for example, the node name, node ID, or another suitable reference or identifier) of the existing node.

In particular embodiments, as a user types or otherwise enters text into a form used to add content or make declarations in various sections of the user’s profile page or other page, the typeahead process may work in conjunction with one or more frontend (client-side) and/or backend (server-side) typeahead processes (hereinafter referred to simply as “typeahead process”) executing at (or within) the social-networking system 160 (e.g., within servers 162), to interactively and virtually instantaneously (as appearing to the user) attempt to auto-populate the form with a term or terms corresponding to names of existing social-graph entities, or terms associated with existing social-graph entities, determined to be the most relevant or best match to the characters of text entered by the user as the user enters the characters of text. Utilizing the social-graph information in a social-graph database or information extracted and indexed from the social-graph database, including information associated with nodes and edges, the typeahead processes, in conjunction with the information from the social-graph database, as well as potentially in conjunction with various others processes, applications, or databases located within or executing within social-networking system 160, are able to predict a user’s intended declaration with a high degree of precision. However, social-networking system 160 also provides user’s with the freedom to enter any declaration they wish enabling users to express themselves freely.

In particular embodiments, as a user enters text characters into a form box or other field, the typeahead processes may attempt to identify existing social-graph elements (e.g., user nodes 202, concept nodes 204, or edges 206) that match the string of characters entered in the user’s declaration as the user is entering the characters. In particular embodiments, as the user enters characters into a form box, the typeahead process may read the string of entered textual characters. As each keystroke is made, the frontend-typeahead process may transmit the entered character string as a request (or call) to the backend-typeahead process executing within social-networking system 160. In particular embodiments, the typeahead processes may communicate via AJAX (Asynchronous JavaScript and XML) or other suitable techniques, and particularly, asynchronous techniques. In one particular embodiment, the request is, or comprises, an XMLHttpRequest (XHR) enabling quick and dynamic sending and fetching of results. In particular embodiments, the typeahead process also transmits before, after, or with the request a section identifier (section ID) that identifies the particular section of the particular page in which the user is making the declaration. In particular embodiments, a user ID parameter may also be sent, but this may be unnecessary in some embodiments, as the user is already “known” based on he or she logging into social-networking system 160.

In particular embodiments, the typeahead process may use one or more matching algorithms to attempt to identify matching social-graph elements. In particular embodiments, when a match or matches are found, the typeahead process may transmit a response (which may utilize AJAX or other suitable techniques) to the user’s client system 130 that may include, for example, the names (name strings) of the matching social-graph elements as well as, potentially, other meta-data associated with the matching social-graph elements. As

an example and not by way of limitation, if a user entering the characters “pok” into a query field, the typeahead process may display a drop-down menu that displays names of matching existing profile pages and respective user nodes **202** or concept nodes **204** (e.g., a profile page named or devoted to “poker”), which the user can then click on or otherwise select thereby confirming the desire to declare the matched user or concept name corresponding to the selected node. As another example and not by way of limitation, upon clicking “poker,” the typeahead process may auto-populate, or causes the web browser **132** to auto-populate, the query field with the declaration “poker”. In particular embodiments, the typeahead process may simply auto-populate the field with the name or other identifier of the top-ranked match rather than display a drop-down menu. The user may then confirm the auto-populated declaration simply by keying “enter” on his or her keyboard or by clicking on the auto-populated declaration.

More information on typeahead processes may be found in U.S. patent application Ser. No. 12/763,162, filed 19 Apr. 2010, and U.S. patent application Ser. No. 13/556,072, filed 23 Jul. 2012, which are incorporated by reference.

Structured Search Queries

FIG. 3 illustrates an example webpage of an online social network. In particular embodiments, a user may submit a query to the social-network system **160** by inputting a text query into query field **350**. A user of an online social network may search for information relating to a specific subject matter (e.g., users, concepts, external content or resource) by providing a short phrase describing the subject matter, often referred to as a “search query,” to a search engine. The query may be an unstructured text query and may comprise one or more text strings or one or more n-grams. In general, a user may input any character string into query field **350** to search for content on the social-networking system **160** that matches the text query. The social-networking system **160** may then search a data store **164** (or, more particularly, a social-graph database) to identify content matching the query. The search engine may conduct a search based on the query phrase using various search algorithms and generate search results that identify resources or content (e.g., user-profile pages, content-profile pages, or external resources) that are most likely to be related to the search query. To conduct a search, a user may input or transmit a search query to the search engine. In response, the search engine may identify one or more resources that are likely to be related to the search query, which may collectively be referred to as a “search result” identified for the search query. The identified content may include, for example, social-graph entities (i.e., user nodes **202**, concept nodes **204**, edges **206**), profile pages, external webpages, or any combination thereof. The social-networking system **160** may then generate a search results webpage with search results corresponding to the identified content. The search results may be presented to the user, often in the form of a list of links on search-results webpage, each link being associated with a different webpage that contains some of the identified resources or content. In particular embodiments, each link in the search results may be in the form of a Uniform Resource Locator (URL) that specifies where the corresponding webpage is located and the mechanism for retrieving it. The social-networking system **160** may then transmit the search results webpage to the user’s web browser **132** on the user’s client system **130**. The user may then click on the URL links or otherwise select the content from the search results webpage to access the content from the social-networking system **160** or from an external system, as appropriate. The resources may be ranked and presented to the user according to their relative degrees of relevance to the search

query. The search results may also be ranked and presented to the user according to their relative degree of relevance to the user. In other words, the search results may be personalized for the querying user based on, for example, social-graph information, user information, search or browsing history of the user, or other suitable information related to the user. In particular embodiments, ranking of the resources may be determined by a ranking algorithm implemented by the search engine. As an example and not by way of limitation, resources that are more relevant to the search query or to the user may be ranked higher than the resources that are less relevant to the search query or the user. In particular embodiments, the search engine may limit its search to resources and content on the online social network. However, in particular embodiments, the search engine may also search for resources or contents on other sources, such as a third-party system **170**, the internet or World Wide Web, or other suitable sources. Although this disclosure describes querying the social-networking system **160** in a particular manner, this disclosure contemplates querying the social-networking system **160** in any suitable manner.

In particular embodiments, the typeahead processes described herein may be applied to search queries entered by a user. As an example and not by way of limitation, as a user enters text characters into a search field, a typeahead process may attempt to identify one or more user nodes **202**, concept nodes **204**, or edges **206** that match the string of characters entered search field as the user is entering the characters. As the typeahead process receives requests or calls including a string or n-gram from the text query, the typeahead process may perform or causes to be performed a search to identify existing social-graph elements (i.e., user nodes **202**, concept nodes **204**, edges **206**) having respective names, types, categories, or other identifiers matching the entered text. The typeahead process may use one or more matching algorithms to attempt to identify matching nodes or edges. When a match or matches are found, the typeahead process may transmit a response to the user’s client system **130** that may include, for example, the names (name strings) of the matching nodes as well as, potentially, other metadata associated with the matching nodes. The typeahead process may then display a drop-down menu **300** that displays names of matching existing profile pages and respective user nodes **202** or concept nodes **204**, and displays names of matching edges **206** that may connect to the matching user nodes **202** or concept nodes **204**, which the user can then click on or otherwise select thereby confirming the desire to search for the matched user or concept name corresponding to the selected node, or to search for users or concepts connected to the matched users or concepts by the matching edges. Alternatively, the typeahead process may simply auto-populate the form with the name or other identifier of the top-ranked match rather than display a drop-down menu **300**. The user may then confirm the auto-populated declaration simply by keying “enter” on a keyboard or by clicking on the auto-populated declaration. Upon user confirmation of the matching nodes and edges, the typeahead process may transmit a request that informs the social-networking system **160** of the user’s confirmation of a query containing the matching social-graph elements. In response to the request transmitted, the social-networking system **160** may automatically (or alternately based on an instruction in the request) call or otherwise search a social-graph database for the matching social-graph elements, or for social-graph elements connected to the matching social-graph elements as appropriate. Although this disclosure describes applying the typeahead processes to search queries in a particular manner,

this disclosure contemplates applying the typeahead processes to search queries in any suitable manner.

Parsing Queries Using Context-Free Grammar Models

FIGS. 4A-4B illustrate example queries of the social network. In particular embodiments, the social-networking system 160 may generate one or more structured queries comprising query tokens that correspond to one or more identified social-graph elements in response to a text query received from a first user (i.e., the querying user). FIGS. 4A-4B illustrate various example text queries in query field 350 and various structured queries generated in response in drop-down menus 300. By providing suggested structured queries in response to a user's text query, the social-networking system 160 may provide a powerful way for users of the online social network to search for elements represented in the social graph 200 based on their social-graph attributes and their relation to various social-graph elements. Structured queries may allow a querying user to search for content that is connected to particular users or concepts in the social graph 200 by particular edge types. As an example and not by way of limitation, the social-networking system 160 may receive a substantially unstructured text query from a first user. In response, the social-networking system 160 (via, for example, a server-side element detection process) may access the social graph 200 and then parse the text query to identify social-graph elements that corresponded to n-grams from the text query. The social-networking system 160 may identify these corresponding social-graph elements by determining a probability for each n-gram that it corresponds to a particular social-graph element. The social-networking system 160 may then access a grammar model, such as a context-free grammar model. The identified social-graph elements may be used as terminal tokens ("query tokens") in the grammars, and each grammar may then be scored. Grammars with a score greater than a threshold score may be used to generate structured queries that include query tokens referencing the identified social-graph elements. The structured queries may then be transmitted to the first user and displayed in a drop-down menu 300 (via, for example, a client-side typeahead process), where the first user can then select an appropriate query to search for the desired content. Some of the advantages of using the structured queries described herein include finding users of the online social networking based upon limited information, bringing together virtual indexes of content from the online social network based on the relation of that content to various social-graph elements, or finding content related to you and/or your friends. Although this disclosure describes and FIGS. 4A-4B illustrate generating particular structured queries in a particular manner, this disclosure contemplates generating any suitable structured queries in any suitable manner.

In particular embodiments, the social-networking system 160 may receive from a querying/first user (corresponding to a first user node 202) a substantially unstructured text query. As an example and not by way of limitation, a first user may want to search for other users who: (1) are first-degree friends of the first user; and (2) are associated with Stanford University (i.e., the user nodes 202 are connected by an edge 206 to the concept node 204 corresponding to the school "Stanford"). The first user may then enter a text query "friends stanford" into query field 350, as illustrated in FIGS. 4A-4B. As the first user enters this text query into query field 350, the social-networking system 160 may provide various suggested structured queries, as illustrated in drop-down menus 300. As used herein, a substantially unstructured text query refers to a simple text string inputted by a user. The text query may, of course, be structured with respect to standard lan-

guage/grammar rules (e.g. English language grammar). However, the text query will ordinarily be unstructured with respect to social-graph elements. In other words, a simply text query will not ordinarily include embedded references to particular social-graph elements. Thus, as used herein, a structured query refers to a query that contains references to particular social-graph elements, allowing the search engine to search based on the identified elements. Although this disclosure describes receiving particular queries in a particular manner, this disclosure contemplates receiving any suitable queries in any suitable manner.

In particular embodiments, social-networking system 160 may parse the substantially unstructured text query (also simply referred to as a search query) received from the first user (i.e., the querying user) to identify one or more n-grams. In general, an n-gram is a contiguous sequence of n items from a given sequence of text or speech. The items may be characters, phonemes, syllables, letters, words, base pairs, prefixes, or other identifiable items from the sequence of text or speech. The n-gram may comprise one or more characters of text (letters, numbers, punctuation, etc.) entered by the querying user. An n-gram of size one can be referred to as a "unigram," of size two can be referred to as a "bigram" or "digram," of size three can be referred to as a "trigram," and so on. Each n-gram may include one or more parts from the text query received from the querying user. In particular embodiments, each n-gram may comprise a character string (e.g., one or more characters of text) entered by the first user. As an example and not by way of limitation, the social-networking system 160 may parse the text query "friends stanford" to identify the following n-grams: friends; stanford; friends stanford. As another example and not by way of limitation, the social-networking system 160 may parse the text query "friends in palo alto" to identify the following n-grams: friends; in; palo; alto; friends in; in palo; palo alto; friend in palo; in palo also; friends in palo alto. In particular embodiments, each n-gram may comprise a contiguous sequence of n items from the text query. Although this disclosure describes parsing particular queries in a particular manner, this disclosure contemplates parsing any suitable queries in any suitable manner.

In particular embodiments, social-networking system 160 may determine or calculate, for each n-gram identified in the text query, a score that the n-gram corresponds to a social-graph element. The score may be, for example, a confidence score, a probability, a quality, a ranking, another suitable type of score, or any combination thereof. As an example and not by way of limitation, the social-networking system 160 may determine a probability score (also referred to simply as a "probability") that the n-gram corresponds to a social-graph element, such as a user node 202, a concept node 204, or an edge 206 of social graph 200. The probability score may indicate the level of similarity or relevance between the n-gram and a particular social-graph element. There may be many different ways to calculate the probability. The present disclosure contemplates any suitable method to calculate a probability score for an n-gram identified in a search query. In particular embodiments, the social-networking system 160 may determine a probability, p, that an n-gram corresponds to a particular social-graph element. The probability, p, may be calculated as the probability of corresponding to a particular social-graph element, k, given a particular search query, X. In other words, the probability may be calculated as $p=(k|X)$. As an example and not by way of limitation, a probability that an n-gram corresponds to a social-graph element may be calculated as an probability score denoted as $p_{i,j,k}$. The input may be a text query $X=(x_1, x_2, \dots, x_N)$, and a set of classes. For each

(i;j) and a class k , the social-networking system **160** may compute $p_{i,j,k} = p(\text{class}(x_{i,j}) = k | X)$. As an example and not by way of limitation, the n-gram “stanford” could be scored with respect to the following social-graph elements as follows: school “Stanford University”=0.7; location “Stanford, Calif.”=0.2; user “Allen Stanford”=0.1. As another example and not by way of limitation, the n-gram “friends” could be scored with respect to the following social-graph elements as follows: user “friends”=0.9; television show “Friends”=0.1. In particular embodiments, the social-networking system **160** may use a forward-backward algorithm to determine the probability that a particular n-gram corresponds to a particular social-graph element. For a given n-gram within a text query, the social-networking system **160** may use both the preceding and succeeding n-grams to determine which particular social-graph elements correspond to the given n-gram. Although this disclosure describes determining whether n-grams correspond to social-graph elements in a particular manner, this disclosure contemplates determining whether n-grams correspond to social-graph elements in any suitable manner. Moreover, although this disclosure describes determining whether an n-gram corresponds to a social-graph element using a particular type of score, this disclosure contemplates determining whether an n-gram corresponds to a social-graph element using any suitable type of score.

In particular embodiments, social-networking system **160** may identify one or more edges **206** having a probability greater than an edge-threshold probability. Each of the identified edges **206** may correspond to at least one of the n-grams. As an example and not by way of limitation, the n-gram may only be identified as corresponding to an edge k if $p_{i,j,k} > p_{\text{edge-threshold}}$. Furthermore, each of the identified edges **206** may be connected to at least one of the identified nodes. In other words, the social-networking system **160** may only identify edges **206** or edge-types that are connected to user nodes **202** or concept nodes **204** that have previously been identified as corresponding to a particular n-gram. Edges **206** or edge-types that are not connected to any previously identified node are typically unlikely to correspond to a particular n-gram in a search query. By filtering out or ignoring these edges **206** and edge-types, the social-networking system **160** may more efficiently search the social graph for relevant social-graph elements. As an example and not by way of limitation, referencing FIG. 2, for a text query containing “went to Stanford,” where an identified concept node **204** is the school “Stanford,” the social-networking system **160** may identify the edges **206** corresponding to “worked at” and the edges **206** corresponding to “attended,” both of which are connected to the concept node **204** for “Stanford.” Thus, the n-gram “went to” may be identified as corresponding to these edges **206**. However, for the same text query, the social-networking system **160** may not identify the edges **206** corresponding to “like” or “fan” in the social graph **200** because the “Stanford” node does not have any such edges connected to it. Although this disclosure describes identifying edges **206** that corresponding to n-grams in a particular manner, this disclosure contemplates identifying edges **206** that corresponding to n-grams in any suitable manner.

In particular embodiments, social-networking system **160** may identify one or more user nodes **202** or concept nodes **204** having a probability greater than a node-threshold probability. Each of the identified nodes may correspond to at least one of the n-grams. As an example and not by way of limitation, the n-gram may only be identified as corresponding to a node k if $p_{i,j,k} > p_{\text{node-threshold}}$. Furthermore, each of the identified user nodes **202** or concept nodes **204** may be connected to at least one of the identified edges **206**. In other words, the

social-networking system **160** may only identify nodes or nodes-types that are connected to edges **206** that have previously been identified as corresponding to a particular n-gram. Nodes or node-types that are not connected to any previously identified edges **206** are typically unlikely to correspond to a particular n-gram in a search query. By filtering out or ignoring these nodes and node-types, the social-networking system **160** may more efficiently search the social graph for relevant social-graph elements. As an example and not by way of limitation, for a text query containing “worked at Apple,” where an identified edge **206** is “worked at,” the social-networking system **160** may identify the concept node **204** corresponding to the company APPLE, INC., which may have multiple edges **206** of “worked at” connected to it. However, for the same text query, the social-networking system **160** may not identify the concept node **204** corresponding to the fruit-type “apple,” which may have multiple “like” or “fan” edges connected to it, but no “worked at” edge connections. In particular embodiments, the node-threshold probability may differ for user nodes **202** and concept nodes **204**. The n-gram may be identified as corresponding to a user node **302** k_{user} if $p_{i,j,k} > p_{\text{user-node-threshold}}$ while the n-gram may be identified as corresponding to a concept node **304** k_{concept} if $p_{i,j,k} > p_{\text{concept-node-threshold}}$. In particular embodiments, the social-networking system **160** may only identify nodes that are within a threshold degree of separation of the user node **202** corresponding to the first user (i.e., the querying user). The threshold degree of separation may be, for example, one, two, three, or all. Although this disclosure describes identifying nodes that corresponding to n-grams in a particular manner, this disclosure contemplates identifying nodes that corresponding to n-grams in any suitable manner.

In particular embodiments, the social-networking system **160** may access a context-free grammar model comprising a plurality of grammars. Each grammar of the grammar model may comprise one or more non-terminal symbols that may be replaced by query tokens. A grammar model is a set of formation rules for strings in a formal language. To generate a string in the language, one begins with a string consisting of only a single start symbol. The production rules are then applied in any order, until a string that contains neither the start symbol nor designated non-terminal symbols is produced. In a context-free grammar, the production of each non-terminal symbol of the grammar is independent of what is produced by other non-terminal symbols of the grammar. The non-terminal symbols may be replaced with terminal symbols (i.e., terminal tokens or query tokens). Some of the query tokens may correspond to identified nodes or identified edges, as described previously. A string generated by the grammar may then be used as a structured query containing references to the identified nodes or identified edges. A context-free grammar is a grammar in which the left-hand side of each production rule consists of only a single non-terminal symbol. A probabilistic context-free grammar is a tuple $\langle E, N, S, P \rangle$, where the disjoint sets Σ and N specify the terminal and non-terminal symbols, respectively, with $S \in N$ being the start symbol. P is the set of productions, which take the form $E \rightarrow \xi(p)$, with $E \in N$, $\xi \in (\Sigma \cup N)^+$, and $p = \text{Pr}(E \rightarrow \xi)$, the probability that E will be expanded into the string ξ . The sum of probabilities p over all expansions of a given non-terminal E must be one. Although this disclosure describes accessing particular grammars, this disclosure contemplates any suitable grammars.

In particular embodiments, the social-networking system **160** may identify one or more grammars having query tokens corresponding to the previously identified nodes and edges. In other words, if an identified node or identified edge may be

used as a query token in a particular grammar, that grammar may be identified by the social-networking system 160 as a possible grammar to use for generating a structured query. This is effectively a type of bottom-up parsing, where the possible query tokens are used to determine the applicable grammar to apply to the query. As an example and not by way of limitation, an example grammar may be: [user] [user-filter] [school]. The non-terminal symbols [user], [user-filter], and [school] could then be determined based n-grams in the received text query. For the text query “friends stanford”, this query could be parsed by using the grammar as, for example, “[friends] [who go to] [Stanford University]” or “[friends] [who work at] [Stanford University]”. As another example and not by way of limitation, an example grammar may be [user] [user-filter] [location]. For the text query “friends stanford”, this query could be parsed by using the grammar as, for example, “[friends][who live in][Stanford, Calif.]”. In both the example cases above, if the n-grams of the received text query could be used as query tokens in the grammars, then these grammars may be identified by the social-networking system 160. Similarly, if the received text query comprises n-grams that could not be used as query tokens in the grammar, that grammar may not be identified. Although this disclosure describes identifying particular grammars in a particular manner, this disclosure contemplates identifying any suitable grammars in any suitable manner.

In particular embodiments, the social-networking system 160 may determine a score for each identified grammar. The score may be, for example, a confidence score, a probability, a quality, a ranking, another suitable type of score, or any combination thereof. The score may be based on the individual scores or probabilities associated with the query tokens of the grammar. A grammar may have a higher relative score if it uses query tokens with relatively higher individual scores. As an example and not by way of limitation, continuing with the prior examples, the n-gram “stanford” could be scored with respect to the following social-graph elements as follows: school “Stanford University”=0.7; location “Stanford, Calif.”=0.2; user “Allen Stanford”=0.1. The n-gram “friends” could be scored with respect to the following social-graph elements as follows: user “friends”=0.9; television show “Friends”=0.1. Thus, the grammar [user][user-filter] [school] may have a relatively high score if it uses the query tokens for the user “friends” and the school “Stanford University” (generating, for example, the string “friends who go to Stanford University”), both of which have relatively high individual scores. In contrast, the grammar [user][user-filter] [user] may have relatively low score if it uses the query tokens for the user “friends” and the user “Allen Stanford” (generating, for example, the string “friends of Allen Stanford”), since the latter query token has a relatively low individual score. Although this disclosure describes determining particular scores for particular grammars in a particular manner, this disclosure contemplates determining any suitable scores for any suitable grammars in any suitable manner.

In particular embodiments, the social-networking system 160 may determine the score for an identified grammar based on the relevance of the social-graph elements corresponding to the query tokens of the grammar to the querying user (i.e., the first user, corresponding to a first user node 202). User nodes 202 and concept nodes 204 that are connected to the first user node 202 directly by an edge 206 may be considered relevant to the first user. Thus, grammars comprising query tokens corresponding to these relevant nodes and edges may be considered more relevant to the querying user. As an example and not by way of limitation, a concept node 204 connected by an edge 206 to a first user node 202 may be

considered relevant to the first user node 202. As used herein, when referencing a social graph 200 the term “connected” means a path exists within the social graph 200 between two nodes, wherein the path may comprise one or more edges 206 and zero or more intermediary nodes. In particular embodiments, nodes that are connected to the first user node 202 via one or more intervening nodes (and therefore two or more edges 206) may also be considered relevant to the first user. Furthermore, in particular embodiments, the closer the second node is to the first user node, the more relevant the second node may be considered to the first user node. That is, the fewer edges 206 separating the first user node 202 from a particular user node 202 or concept node 204 (i.e., the fewer degrees of separation), the more relevant that user node 202 or concept node 204 may be considered to the first user. As an example and not by way of limitation, as illustrated in FIG. 2, the concept node 204 corresponding to the school “Stanford” is connected to the user node 202 corresponding to User “C,” and thus the concept “Stanford” may be considered relevant to User “C.” As another example and not by way of limitation, the user node 202 corresponding to User “A” is connected to the user node 202 corresponding to User “C” via one intermediate node and two edges 206 (i.e., the intermediated user node 202 corresponding to User “B”), and thus User “A” may be considered relevant to User “C,” but because the user node 202 for User “A” is a second-degree connection with respect to User “C,” that particular concept node 204 may be considered less relevant than a user node 202 that is connected to the user node for User “C” by a single edge 206, such as, for example, the user node 202 corresponding to User “B.” As yet another example and not by way of limitation, the concept node for “Online Poker” (which is an online multiplayer game) is not connected to the user node for User “C” by any pathway in social graph 200, and thus the concept “Online Poker” may not be considered relevant to User “C.” In particular embodiments, a second node may only be considered relevant to the first user if the second node is within a threshold degree of separation of the first user node 202. As an example and not by way of limitation, if the threshold degree of separation is three, then the user node 202 corresponding to User “D” may be considered relevant to the concept node 204 corresponding to the recipe “Chicken Parmesan,” which are within three degrees of each other on social graph 200 illustrated in FIG. 2. However, continuing with this example, the concept node 204 corresponding to the application “All About Recipes” would not be considered relevant to the user node 202 corresponding to User “D” because these nodes are four degrees apart in the social graph 200. Although this disclosure describes determining whether particular social-graph elements (and thus their corresponding query tokens) are relevant to each other in a particular manner, this disclosure contemplates determining whether any suitable social-graph elements are relevant to each other in any suitable manner. Moreover, although this disclosure describes determining whether particular query tokens corresponding to user nodes 202 and concept nodes 204 are relevant to a querying user, this disclosure contemplates similarly determining whether any suitable query token (and thus any suitable node) is relevant to any other suitable user.

In particular embodiments, the social-networking system 160 may determine the score for an identified grammar based on social-graph information corresponding to the query tokens of the grammar. As an example and not by way of limitation, when determining a probability, p, that an n-gram corresponds to a particular social-graph element, the calculation of the probability may also factor in social-graph information. Thus, the probability of corresponding to a particular social-

graph element, k , given a particular search query, X , and social-graph information, G , may be calculated as $p=(k|X, G)$. The individual probabilities for the identified nodes and edges may then be used to determine the score for a grammar using those social-graph elements as query tokens. In particular embodiments, the score for an identified grammar may be based on the degree of separation between the first user node **202** and the particular social-graph element used as a query token in the grammar. Grammars with query tokens corresponding to social-graph elements that are closer in the social graph **200** to the querying user (i.e., fewer degrees of separation between the element and the first user node **202**) than a social-graph element that is further from the user (i.e., more degrees of separation). As an example and not by way of limitation, referencing FIG. 2, if user “B” inputs a text query of “chicken,” a grammar with a query token corresponding to the concept node **204** for the recipe “Chicken Parmesan,” which is connected to user “B” by an edge **206**, may have a relatively higher score than a grammar with a query token corresponding to other nodes associated with the n-gram chicken (e.g., concept nodes **204** corresponding to “chicken nuggets,” or “funky chicken dance”) that are not connected to user “B” in the social graph **200**. In particular embodiments, the score for an identified grammar may be based on the identified edges **206** corresponding to the query tokens of the grammar. If the social-networking system **160** has already identified one or more edges that correspond to n-grams in a received text query, those identified edges may then be considered when determining the score for a particular parsing of the text query by the grammar. If a particular grammar comprises query tokens that correspond to both identified nodes and identified edges, if the identified nodes are not actually connected to any of the identified edges, that particular grammar may be assigned a zero or null score. In particular embodiments, the score for an identified grammar may be based on the number of edges **206** connected to the node corresponding to a query token of the grammar. Grammars comprising query tokens that corresponding to nodes with more connecting edges **206** may be more popular and more likely to be a target of a search query. As an example and not by way of limitation, if the concept node **204** for “Stanford, Calif.” is only connected by five edges while the concept node **204** for “Stanford University” is connected by five-thousand edges, when determining the score for grammars containing query tokens corresponding to either of these nodes, the social-networking system **160** may determine that the grammar referencing the concept node **204** for “Stanford University” has a relatively higher score than a grammar referencing the concept node **204** for “Stanford, Calif.” because of the greater number of edges connected to the former concept node **204**. In particular embodiments, the score for an identified grammar may be based on the search history associate with the first user (i.e., the querying user). Grammars with query tokens corresponding to nodes that the first user has previously accessed, or are relevant to the nodes the first user has previously accessed, may be more likely to be the target of the first user’s search query. Thus, these grammars may be given a higher score. As an example and not by way of limitation, if first user has previously visited the “Stanford University” profile page but has never visited the “Stanford, Calif.” profile page, when determining the score for grammars with query tokens corresponding to these concepts, the social-networking system **160** may determine that the concept node **204** for “Stanford University” has a relatively high score, and thus the grammar using the corresponding query token, because the querying user has previously accessed the concept node **204** for the school. As another example and not

by way of limitation, if the first user has previously visited the concept-profile page for the television show “Friends,” when determining the score for the grammar with the query token corresponding to that concept, the social-networking system **160** may determine that the concept node **204** corresponding to the television show “Friends” has a relatively high score, and thus the grammar using the corresponding query token, because the querying user has previously accessed the concept node **204** for that television show. Although this disclosure describes determining scores for particular grammars based on particular social-graph information in a particular manner, this disclosure contemplates determining scores for any suitable grammars based on any suitable social-graph information in any suitable manner.

In particular embodiments, social-networking system **160** may select one or more grammars having a score greater than a grammar-threshold score. Each of the selected grammars may contain query tokens that correspond to at least one of the identified nodes or identified edges (which correspond to n-grams of the received text query). In particular embodiments, the grammars may be ranked based on their determined scores, and only grammars within a threshold rank may be selected (e.g., top seven). Although this disclosure describes selecting grammars in a particular manner, this disclosure contemplates selecting grammars in any suitable manner.

In particular embodiments, social-networking system **160** may generate one or more structured queries corresponding to an identified grammar having a score greater than a grammar-threshold score. Each structure query may be based on a string generated by the corresponding identified grammar. As an example and not by way of limitation, in response to the text query “friends stanford”, the grammar [user] [user-filter] [school] may generate a string “friends who go to Stanford University”, where the non-terminal tokens [user], [user-filter], [school] of the grammar have been replaced by the terminal tokens [friends], [who go to], and [Stanford University], respectively, to generate the string. Each structured query may comprise query tokens corresponding to the corresponding identified grammar, where these query tokens correspond to one or more of the identified edges **206** and one or more of the identified nodes. Generating structured queries is described more below.

FIG. 5 illustrates an example method **500** for using a context-free grammar model to generate structured search queries. The method may begin at step **510**, where the social-networking system **160** may access a social graph **200** comprising a plurality of nodes and a plurality of edges **206** connecting the nodes. The nodes may comprise a first user node **202** and a plurality of second nodes (one or more user nodes **202**, concepts nodes **204**, or any combination thereof). At step **520**, the social-networking system **160** may receive from the first user a substantially unstructured text query. The text query may comprise one or more n-grams. At step **530**, the social-networking system **160** may identify edges and second nodes corresponding to the n-grams. At step **540**, the social-networking system **160** may access a context-free grammar model comprising a plurality of grammars. Each grammar may comprise one or more query tokens. At step **550**, the social-networking system **160** may identify grammars having query tokens corresponding to the identified nodes or identified edges. At step **560**, the social-networking system **160** may determine a score for each identified grammar. This score may be based on a variety of factors. At step **570**, the social-networking system may generate one or more structured queries based on the identified grammars. Each structured query may correspond to an identified grammar

having a score greater than a grammar-threshold score, and may comprise the query tokens of the corresponding identified grammar. The query tokens of the structured query may correspond to at least one of the identified second nodes or identified edges. Particular embodiments may repeat one or more steps of the method of FIG. 5, where appropriate. Although this disclosure describes and illustrates particular steps of the method of FIG. 5 as occurring in a particular order, this disclosure contemplates any suitable steps of the method of FIG. 5 occurring in any suitable order. Moreover, although this disclosure describes and illustrates particular components, devices, or systems carrying out particular steps of the method of FIG. 5, this disclosure contemplates any suitable combination of any suitable components, devices, or systems carrying out any suitable steps of the method of FIG. 5.

Generating Structured Search Queries

In particular embodiments, social-networking system 160 may generate one or more structured queries that each comprise the query tokens of the corresponding grammar, where the query tokens may correspond to one or more of the identified user nodes 202 or one or more of the identified edges 206. The generated structured queries may be based on context-free grammars, as described previously. This type of structured search query may allow the social-networking system 160 to more efficiently search for resources and content related to the online social network (such as, for example, profile pages) by searching for content connected to or otherwise related to the identified user nodes 202 and the identified edges 206. As an example and not by way of limitation, in response to the text query, “show me friends of my girlfriend,” the social-networking system 160 may generate a structured query “Friends of Stephanie,” where “Friends” and “Stephanie” in the structured query are references corresponding to particular social-graph elements. The reference to “Stephanie” would correspond to a particular user node 202, while the reference to “friends” would correspond to “friend” edges 206 connecting that user node 202 to other user nodes 202 (i.e., edges 206 connecting to “Stephanie’s” first-degree friends). When executing this structured query, the social-networking system 160 may identify one or more user nodes 202 connected by “friend” edges 206 to the user node 202 corresponding to “Stephanie.” In particular embodiments, the social-networking system 160 may generate a plurality of structured queries, where the structured queries may comprise references to different identified user nodes 202 or different identified edges 206. As an example and not by way of limitation, in response to the text query, “photos of cat,” the social-networking system 160 may generate a first structured query “Photos of Catey” and a second structured query “Photos of Catherine,” where “Photos” in the structured query is a reference corresponding to a particular social-graph element, and where “Catey” and “Catherine” are references to two different user nodes 202. When executing either of these structured queries, the social-networking system 160 may identify one or more concept nodes 204 corresponding to photos that are connected to the identified user nodes 202 by edges 206. Although this disclosure describes generating particular structured queries in a particular manner, this disclosure contemplates generating any suitable structured queries in any suitable manner.

In particular embodiments, social-networking system 160 may generate one or more structured queries that each comprise query tokens corresponding to the identified concept nodes 204 and one or more of the identified edges 206. This type of structured search query may allow the social-networking system 160 to more efficiently search for resources and

content related to the online social network (such as, for example, profile pages) by search for content connected to or otherwise related to the identified concept nodes 204 and the identified edges 206. As an example and not by way of limitation, in response to the text query, “friends who like facebook,” the social-networking system 160 may generate a structured query “Friends who like Facebook,” where “Friends,” like,” and “Facebook” in the structured query are query tokens corresponding to particular social-graph elements as described previously (i.e., a “friend” edge 206, a “like” edge 206, and a “Facebook” concept node 204). In particular embodiments, the social-networking system 160 may generate a plurality of structured queries, where the structured queries may comprise references to different identified concept nodes 204 or different identified edges 206. As an example and not by way of limitation, continuing with the previous example, in addition to the structured query “Friends who like Facebook,” the social-networking system 160 may also generate a structured query “Friends who like Facebook Culinary Team,” where “Facebook Culinary Team” in the structured query is a query token corresponding to yet another social-graph element. In particular embodiments, social-networking system 160 may rank the generated structured queries. The structured queries may be ranked based on a variety of factors. Although this disclosure describes generating particular structured queries in a particular manner, this disclosure contemplates generating any suitable structured queries in any suitable manner.

In particular embodiments, social-networking system 160 may transmit one or more of the structured queries to the first user (i.e., the querying user). As an example and not by way of limitation, after the structured queries are generated, the social-networking system 160 may transmit one or more of the structured queries as a response (which may utilize AJAX or other suitable techniques) to the user’s client system 130 that may include, for example, the names (name strings) of the referenced social-graph elements, other query limitations (e.g., Boolean operators, etc.), as well as, potentially, other metadata associated with the referenced social-graph elements. The web browser 132 on the querying user’s client system 130 may display the transmitted structured queries in a drop-down menu 300, as illustrated in FIGS. 4A-4B. In particular embodiments, the transmitted queries may be presented to the querying user in a ranked order, such as, for example, based on a rank previously determined as described above. Structured queries with better rankings may be presented in a more prominent position. Furthermore, in particular embodiments, only structured queries above a threshold rank may be transmitted or displayed to the querying user. As an example and not by way of limitation, as illustrated in FIGS. 4A-4B, the structured queries may be presented to the querying user in a drop-down menu 300 where higher ranked structured queries may be presented at the top of the menu, with lower ranked structured queries presented in descending order down the menu. In the examples illustrated in FIGS. 4A-4B, only the seven highest ranked queries are transmitted and displayed to the user. In particular embodiments, one or more references in a structured query may be highlighted in order to indicate its correspondence to a particular social-graph element. As an example and not by way of limitation, as illustrated in FIGS. 4A-4B, the references to “Stanford University” and “Stanford, Calif.” may be highlighted in the structured queries to indicate that it corresponds to a particular concept node 204. Similarly, the references to “Friends,” “like,” “work at,” and “go to” in the structured queries presented in drop-down menu 300 could also be highlighted to indicate that they correspond to particular edges 206.

Although this disclosure describes transmitting particular structured queries in a particular manner, this disclosure contemplates transmitting any suitable structured queries in any suitable manner.

In particular embodiments, social-networking system **160** may receive from the first user (i.e., the querying user) a selection of one of the structured queries. As an example and not by way of limitation, the web browser **132** on the querying user's client system **130** may display the transmitted structured queries in a drop-down menu **300**, as illustrated in FIGS. 4A-4B, which the user may then click on or otherwise select (e.g., by simply keying "enter" on his keyboard) to indicate the particular structured query the user wants the social-networking system **160** to execute. Upon selecting the particular structured query, the user's client system **130** may call or otherwise instruct to the social-networking system **160** to execute the selected structured query. Although this disclosure describes receiving selections of particular structured queries in a particular manner, this disclosure contemplates receiving selections of any suitable structured queries in any suitable manner.

More information on structured search queries may be found in U.S. patent application Ser. No. 13/556,072, filed 23 Jul. 2012, which is incorporated by reference.

Systems and Methods

FIG. 6 illustrates an example computer system **600**. In particular embodiments, one or more computer systems **600** perform one or more steps of one or more methods described or illustrated herein. In particular embodiments, one or more computer systems **600** provide functionality described or illustrated herein. In particular embodiments, software running on one or more computer systems **600** performs one or more steps of one or more methods described or illustrated herein or provides functionality described or illustrated herein. Particular embodiments include one or more portions of one or more computer systems **600**. Herein, reference to a computer system may encompass a computing device, where appropriate. Moreover, reference to a computer system may encompass one or more computer systems, where appropriate.

This disclosure contemplates any suitable number of computer systems **600**. This disclosure contemplates computer system **600** taking any suitable physical form. As example and not by way of limitation, computer system **600** may be an embedded computer system, a system-on-chip (SOC), a single-board computer system (SBC) (such as, for example, a computer-on-module (COM) or system-on-module (SOM)), a desktop computer system, a laptop or notebook computer system, an interactive kiosk, a mainframe, a mesh of computer systems, a mobile telephone, a personal digital assistant (PDA), a server, a tablet computer system, or a combination of two or more of these. Where appropriate, computer system **600** may include one or more computer systems **600**; be unitary or distributed; span multiple locations; span multiple machines; span multiple data centers; or reside in a cloud, which may include one or more cloud components in one or more networks. Where appropriate, one or more computer systems **600** may perform without substantial spatial or temporal limitation one or more steps of one or more methods described or illustrated herein. As an example and not by way of limitation, one or more computer systems **600** may perform in real time or in batch mode one or more steps of one or more methods described or illustrated herein. One or more computer systems **600** may perform at different times or at different locations one or more steps of one or more methods described or illustrated herein, where appropriate.

In particular embodiments, computer system **600** includes a processor **602**, memory **604**, storage **606**, an input/output (I/O) interface **608**, a communication interface **610**, and a bus **612**. Although this disclosure describes and illustrates a particular computer system having a particular number of particular components in a particular arrangement, this disclosure contemplates any suitable computer system having any suitable number of any suitable components in any suitable arrangement.

In particular embodiments, processor **602** includes hardware for executing instructions, such as those making up a computer program. As an example and not by way of limitation, to execute instructions, processor **602** may retrieve (or fetch) the instructions from an internal register, an internal cache, memory **604**, or storage **606**; decode and execute them; and then write one or more results to an internal register, an internal cache, memory **604**, or storage **606**. In particular embodiments, processor **602** may include one or more internal caches for data, instructions, or addresses. This disclosure contemplates processor **602** including any suitable number of any suitable internal caches, where appropriate. As an example and not by way of limitation, processor **602** may include one or more instruction caches, one or more data caches, and one or more translation lookaside buffers (TLBs). Instructions in the instruction caches may be copies of instructions in memory **604** or storage **606**, and the instruction caches may speed up retrieval of those instructions by processor **602**. Data in the data caches may be copies of data in memory **604** or storage **606** for instructions executing at processor **602** to operate on; the results of previous instructions executed at processor **602** for access by subsequent instructions executing at processor **602** or for writing to memory **604** or storage **606**; or other suitable data. The data caches may speed up read or write operations by processor **602**. The TLBs may speed up virtual-address translation for processor **602**. In particular embodiments, processor **602** may include one or more internal registers for data, instructions, or addresses. This disclosure contemplates processor **602** including any suitable number of any suitable internal registers, where appropriate. Where appropriate, processor **602** may include one or more arithmetic logic units (ALUs); be a multi-core processor; or include one or more processors **602**. Although this disclosure describes and illustrates a particular processor, this disclosure contemplates any suitable processor.

In particular embodiments, memory **604** includes main memory for storing instructions for processor **602** to execute or data for processor **602** to operate on. As an example and not by way of limitation, computer system **600** may load instructions from storage **606** or another source (such as, for example, another computer system **600**) to memory **604**. Processor **602** may then load the instructions from memory **604** to an internal register or internal cache. To execute the instructions, processor **602** may retrieve the instructions from the internal register or internal cache and decode them. During or after execution of the instructions, processor **602** may write one or more results (which may be intermediate or final results) to the internal register or internal cache. Processor **602** may then write one or more of those results to memory **604**. In particular embodiments, processor **602** executes only instructions in one or more internal registers or internal caches or in memory **604** (as opposed to storage **606** or elsewhere) and operates only on data in one or more internal registers or internal caches or in memory **604** (as opposed to storage **606** or elsewhere). One or more memory buses (which may each include an address bus and a data bus) may couple processor **602** to memory **604**. Bus **612** may include one or

more memory buses, as described below. In particular embodiments, one or more memory management units (MMUs) reside between processor 602 and memory 604 and facilitate accesses to memory 604 requested by processor 602. In particular embodiments, memory 604 includes random access memory (RAM). This RAM may be volatile memory, where appropriate. Where appropriate, this RAM may be dynamic RAM (DRAM) or static RAM (SRAM). Moreover, where appropriate, this RAM may be single-ported or multi-ported RAM. This disclosure contemplates any suitable RAM. Memory 604 may include one or more memories 604, where appropriate. Although this disclosure describes and illustrates particular memory, this disclosure contemplates any suitable memory.

In particular embodiments, storage 606 includes mass storage for data or instructions. As an example and not by way of limitation, storage 606 may include a hard disk drive (HDD), a floppy disk drive, flash memory, an optical disc, a magneto-optical disc, magnetic tape, or a Universal Serial Bus (USB) drive or a combination of two or more of these. Storage 606 may include removable or non-removable (or fixed) media, where appropriate. Storage 606 may be internal or external to computer system 600, where appropriate. In particular embodiments, storage 606 is non-volatile, solid-state memory. In particular embodiments, storage 606 includes read-only memory (ROM). Where appropriate, this ROM may be mask-programmed ROM, programmable ROM (PROM), erasable PROM (EPROM), electrically erasable PROM (EEPROM), electrically alterable ROM (EAROM), or flash memory or a combination of two or more of these. This disclosure contemplates mass storage 606 taking any suitable physical form. Storage 606 may include one or more storage control units facilitating communication between processor 602 and storage 606, where appropriate. Where appropriate, storage 606 may include one or more storages 606. Although this disclosure describes and illustrates particular storage, this disclosure contemplates any suitable storage.

In particular embodiments, I/O interface 608 includes hardware, software, or both providing one or more interfaces for communication between computer system 600 and one or more I/O devices. Computer system 600 may include one or more of these I/O devices, where appropriate. One or more of these I/O devices may enable communication between a person and computer system 600. As an example and not by way of limitation, an I/O device may include a keyboard, keypad, microphone, monitor, mouse, printer, scanner, speaker, still camera, stylus, tablet, touch screen, trackball, video camera, another suitable I/O device or a combination of two or more of these. An I/O device may include one or more sensors. This disclosure contemplates any suitable I/O devices and any suitable I/O interfaces 608 for them. Where appropriate, I/O interface 608 may include one or more device or software drivers enabling processor 602 to drive one or more of these I/O devices. I/O interface 608 may include one or more I/O interfaces 608, where appropriate. Although this disclosure describes and illustrates a particular I/O interface, this disclosure contemplates any suitable I/O interface.

In particular embodiments, communication interface 610 includes hardware, software, or both providing one or more interfaces for communication (such as, for example, packet-based communication) between computer system 600 and one or more other computer systems 600 or one or more networks. As an example and not by way of limitation, communication interface 610 may include a network interface controller (NIC) or network adapter for communicating with an Ethernet or other wire-based network or a wireless NIC

(WNIC) or wireless adapter for communicating with a wireless network, such as a WI-FI network. This disclosure contemplates any suitable network and any suitable communication interface 610 for it. As an example and not by way of limitation, computer system 600 may communicate with an ad hoc network, a personal area network (PAN), a local area network (LAN), a wide area network (WAN), a metropolitan area network (MAN), or one or more portions of the Internet or a combination of two or more of these. One or more portions of one or more of these networks may be wired or wireless. As an example, computer system 600 may communicate with a wireless PAN (WPAN) (such as, for example, a BLUETOOTH WPAN), a WI-FI network, a WI-MAX network, a cellular telephone network (such as, for example, a Global System for Mobile Communications (GSM) network), or other suitable wireless network or a combination of two or more of these. Computer system 600 may include any suitable communication interface 610 for any of these networks, where appropriate. Communication interface 610 may include one or more communication interfaces 610, where appropriate. Although this disclosure describes and illustrates a particular communication interface, this disclosure contemplates any suitable communication interface.

In particular embodiments, bus 612 includes hardware, software, or both coupling components of computer system 600 to each other. As an example and not by way of limitation, bus 612 may include an Accelerated Graphics Port (AGP) or other graphics bus, an Enhanced Industry Standard Architecture (EISA) bus, a front-side bus (FSB), a HYPERTRANSPORT (HT) interconnect, an Industry Standard Architecture (ISA) bus, an INFINIBAND interconnect, a low-pin-count (LPC) bus, a memory bus, a Micro Channel Architecture (MCA) bus, a Peripheral Component Interconnect (PCI) bus, a PCI-Express (PCIe) bus, a serial advanced technology attachment (SATA) bus, a Video Electronics Standards Association local (VLB) bus, or another suitable bus or a combination of two or more of these. Bus 612 may include one or more buses 612, where appropriate. Although this disclosure describes and illustrates a particular bus, this disclosure contemplates any suitable bus or interconnect.

Herein, a computer-readable non-transitory storage medium or media may include one or more semiconductor-based or other integrated circuits (ICs) (such as, for example, field-programmable gate arrays (FPGAs) or application-specific ICs (ASICs)), hard disk drives (HDDs), hybrid hard drives (HHDs), optical discs, optical disc drives (ODDs), magneto-optical discs, magneto-optical drives, floppy diskettes, floppy disk drives (FDDs), magnetic tapes, solid-state drives (SSDs), RAM-drives, SECURE DIGITAL cards or drives, any other suitable computer-readable non-transitory storage media, or any suitable combination of two or more of these, where appropriate. A computer-readable non-transitory storage medium may be volatile, non-volatile, or a combination of volatile and non-volatile, where appropriate.

Herein, “or” is inclusive and not exclusive, unless expressly indicated otherwise or indicated otherwise by context. Therefore, herein, “A or B” means “A, B, or both,” unless expressly indicated otherwise or indicated otherwise by context. Moreover, “and” is both joint and several, unless expressly indicated otherwise or indicated otherwise by context. Therefore, herein, “A and B” means “A and B, jointly or severally,” unless expressly indicated otherwise or indicated otherwise by context.

The scope of this disclosure encompasses all changes, substitutions, variations, alterations, and modifications to the example embodiments described or illustrated herein that a person having ordinary skill in the art would comprehend.

The scope of this disclosure is not limited to the example embodiments described or illustrated herein. Moreover, although this disclosure describes and illustrates respective embodiments herein as including particular components, elements, functions, operations, or steps, any of these embodi- 5 ments may include any combination or permutation of any of the components, elements, functions, operations, or steps described or illustrated anywhere herein that a person having ordinary skill in the art would comprehend. Furthermore, reference in the appended claims to an apparatus or system or a component of an apparatus or system being adapted to, 10 arranged to, capable of, configured to, enabled to, operable to, or operative to perform a particular function encompasses that apparatus, system, component, whether or not it or that particular function is activated, turned on, or unlocked, as long as that apparatus, system, or component is so adapted, arranged, capable, configured, enabled, operable, or opera- 15 tive.

What is claimed is:

1. A method comprising, by a computing device:

accessing a social graph comprising a plurality of nodes and a plurality of edges connecting the nodes, each of the edges between two of the nodes representing a single degree of separation between them, the nodes compris- 20 ing:

a first-user node corresponding to a first user associated with an online social network; and

a plurality of second nodes that each correspond to a concept or a second user associated with the online social network;

receiving, from a client system of the first user, an unstructured text query inputted by the first user, wherein the unstructured text query comprises one or more n-grams; identifying, based on the unstructured text query, one or more edges and one or more second nodes of the social graph, each of the identified edges or identified nodes corresponding to at least one of the n-grams, wherein identifying the one or more edges and one or more second nodes comprises: 30

determining a second score for each n-gram that corresponds to one of the edges or second nodes;

selecting one or more edges having a second score greater than an edge-threshold score, each of the identified edges corresponding to at least one of the n-grams; and 45

selecting one or more second nodes having a second score greater than a node-threshold score, each of the identified second nodes being connected to at least one of the identified edges, each of the identified second nodes corresponding to at least one of the n-grams; 50

accessing a context-free grammar model comprising a plurality of grammars, each grammar comprising one or more query tokens;

identifying one or more grammars, each identified grammar having one or more query tokens corresponding to at least one of the identified second nodes and at least one of the identified edges of the social graph;

determining a first score for each identified grammar;

generating one or more structured queries, each structured query corresponding to an identified grammar having first score greater than a grammar-threshold score wherein the structured query comprises a natural-language string generated by the identified grammar, each structured query comprising the query tokens of the corresponding identified grammar, wherein one or more of the query tokens of the structured query corresponds 65

to at least one of the identified second nodes and at least one of the or identified edges of the social graph; and sending, to the client system of the first user, one or more of the structured queries as suggested queries for display to the first user in response to the unstructured text query inputted by the first user.

2. The method of claim 1, wherein each n-gram comprises one or more characters of text entered by the first user.

3. The method of claim 1, wherein each n-gram comprises a contiguous sequence of n items from the text query. 10

4. The method of claim 1, wherein the second score for each n-gram is a probability that the n-gram corresponds to an edge or a second node.

5. The method of claim 1, wherein determining the first score for each grammar is based on the degree of separation between the first-user node and the identified second nodes corresponding to the query tokens of the grammar. 15

6. The method of claim 1, wherein determining the first score for each grammar is based on the identified edges corresponding to the query tokens of the grammar. 20

7. The method of claim 1, wherein determining the first score for each grammar is based on the number of identified edges connected to the identified second nodes corresponding to the query tokens of the grammar.

8. The method of claim 1, wherein determining the first score for each grammar is based on a search history associated with the first user. 25

9. One or more computer-readable non-transitory storage media embodying software that is operable when executed to: access a social graph comprising a plurality of nodes and a plurality of edges connecting the nodes, each of the edges between two of the nodes representing a single degree of separation between them, the nodes compris- 30 ing:

a first-user node corresponding to a first user associated with an online social network; and

a plurality of second nodes that each correspond to a concept or a second user associated with the online social network;

receive, from a client system of the first user, an unstructured text query inputted by the first user, wherein the unstructured text query comprises one or more n-grams; identify, based on the unstructured text query, one or more edges and one or more second nodes of the social graph, each of the identified edges or identified nodes corresponding to at least one of the n-grams, wherein identifying the one or more edges and one or more second nodes comprises: 35

determine a second score for each n-gram that corresponds to one of the edges or second nodes;

select one or more edges having a second score greater than an edge-threshold score, each of the identified edges corresponding to at least one of the n-grams; and 40

select one or more second nodes having a second score greater than a node-threshold score, each of the identified second nodes being connected to at least one of the identified edges, each of the identified second nodes corresponding to at least one of the n-grams; 45

access a context-free grammar model comprising a plurality of grammars, each grammar comprising one or more query tokens;

identify one or more grammars, each identified grammar having one or more query tokens corresponding to at least one of the identified second nodes and at least one of the identified edges of the social graph;

determine a first score for each identified grammar; 50

29

generate one or more structured queries, each structured query corresponding to an identified grammar having first score greater than a grammar-threshold score wherein the structured query comprises a natural-language string generated by the identified grammar, each structured query comprising the query tokens of the corresponding identified grammar, wherein one or more of the query tokens of the structured query corresponds to at least one of the identified second nodes and at least one of the identified edges of the social graph; and send, to the client system of the first user, one or more of the structured queries as suggested queries for display to the first user in response to the unstructured text query inputted by the first user.

10. The media of claim 9, wherein each n-gram comprises one or more characters of text entered by the first user.

11. The media of claim 9, wherein each n-gram comprises a contiguous sequence of n items from the text query.

12. The media of claim 9, wherein the second score for each n-gram is a probability that the n-gram corresponds to an edge or a second node.

13. The media of claim 9, wherein determining the first score for each grammar is based on the degree of separation between the first-user node and the identified second nodes corresponding to the query tokens of the grammar.

14. The media of claim 9, wherein determining the first score for each grammar is based on the identified edges corresponding to the query tokens of the grammar.

15. The media of claim 9, wherein determining the first score for each grammar is based on the number of identified edges connected to the identified second nodes corresponding to the query tokens of the grammar.

16. A system comprising: one or more processors; and a memory coupled to the processors comprising instructions executable by the processors, the processors operable when executing the instructions to:

- access a social graph comprising a plurality of nodes and a plurality of edges connecting the nodes, each of the edges between two of the nodes representing a single degree of separation between them, the nodes comprising:
 - a first-user node corresponding to a first user associated with an online social network; and
 - a plurality of second nodes that each correspond to a concept or a second user associated with the online social network;

30

receive, from a client system of the first user, an unstructured text query inputted by the first user, wherein the unstructured text query comprises one or more n-grams;

identify, based on the unstructured text query, one or more edges and one or more second nodes of the social graph, each of the identified edges or identified nodes corresponding to at least one of the n-grams, wherein identifying the one or more edges and one or more second nodes comprises:

determine a second score for each n-gram that corresponds to one of the edges or second nodes;

select one or more edges having a second score greater than an edge-threshold score, each of the identified edges corresponding to at least one of the n-grams; and

select one or more second nodes having a second score greater than a node-threshold score, each of the identified second nodes being connected to at least one of the identified edges, each of the identified second nodes corresponding to at least one of the n-grams;

access a context-free grammar model comprising a plurality of grammars, each grammar comprising one or more query tokens;

identify one or more grammars, each identified grammar having one or more query tokens corresponding to at least one of the identified second nodes and at least one of the identified edges of the social graph;

determine a first score for each identified grammar;

generate one or more structured queries, each structured query corresponding to an identified grammar having first score greater than a grammar-threshold score wherein the structured query comprises a natural-language string generated by the identified grammar, each structured query comprising the query tokens of the corresponding identified grammar, wherein one or more of the query tokens of the structured query corresponds to at least one of the identified second nodes and at least one of the identified edges of the social graph; and

send, to the client system of the first user, one or more of the structured queries as suggested queries for display to the first user in response to the unstructured text query inputted by the first user.

* * * * *